Episodic memory: mental time travel or a quantum 'memory wave' function?

Jeremy R. Manning
Dartmouth College, Hanover, NH
jeremy.r.manning@dartmouth.edu

September 30, 2020

Abstract

Where do we “go” when we recollect our past? When remembering a past event, it is intuitive to imagine some part of ourselves mentally “jumping back in time” to when the event occurred. I propose an alternative view, inspired by recent evidence from my lab and others, as well as by re-examining existing models of episodic recall, that suggests that this notion of mentally revisiting any specific moment of our past is at best incomplete and at worst misleading. Instead, I suggest that we retrieve information from our past by mentally casting ourselves back simultaneously to *many* time points from our past, much like a quantum wave function spreading its probability mass over many possible states. This revised conceptual model makes important behavioral and neural predictions about how we retrieve information about our past, and has implications for how we study episodic memory experimentally.

Introduction and overview

How do our brains organize, retrieve, and act upon the ceaseless flow of incoming information and experiences? Tulving (1983) contends that episodic memory is like a sort of “mental time

19 travel,” whereby we send some part of our mental state back along our autobiographical timeline
20 to a specific moment in our past (also see Hassabis et al., 2007; Manning et al., 2011; Schacter et al.,
21 2007; Tulving, 1972). This process has also been described as mentally “jumping back in time” (e.g.,
22 Folkerts et al., 2018; Howard et al., 2012) to when we experienced what we are now remembering.
23 This view has fundamental implications for how we study and build theories about memory: it
24 suggests that we can match up specific moments of recollection with specific moments from the
25 past. These implications extend from how we design episodic memory experiments to how we an-
26 alyze and interpret the behavioral and neural data from those experiments in order to understand
27 the neural and/or cognitive basis of memory retrieval. At its core, characterizing episodic memory
28 as mental time travel makes studying and understanding memory about matching. Specifically,
29 the primary implied questions are about how and when we match up the thoughts (and/or brain
30 activity patterns) we are experiencing at one moment with the thoughts (and/or brain activity pat-
31 terns) we experienced at another previous moment. The implied challenge, then, is to understand
32 the neural (structural) and cognitive (functional) mechanisms that characterize when, how, and
33 why our memory systems allow us to relive those prior experiences. One way of illustrating how
34 this framing misses out on key aspects of memory is to examine how experimentalists and theorists
35 typically conceive of and model a mental construct called *context*.

36 The notion of context was fundamental to Tulving’s characterization of episodic recall, and
37 continues to be a primary feature of most current theories of episodic memory (e.g., for review see
38 Kahana, 2012). Whereas the *content* of a memory provides specific information about the episode
39 itself, contextual components of a memory help to frame the episode within the rememberer’s
40 broader experience. Defining context precisely (e.g., as distinct from content) is an ongoing chal-
41 lenge for our field, but theorists generally describe context as reflecting a mental combination of
42 external cues (where you are, who you are with, background sensory information, etc.) and in-
43 ternal cues (thoughts, feelings, emotions, etc.) that uniquely define each moment (e.g., see review
44 by Manning et al., 2015). Reactivating the mental representation of the context associated with
45 a given episode is what can give us the feeling of some “part of ourselves” re-experiencing the
46 past (Tulving refers to this feeling as *autonoetic awareness*; Tulving, 2002b).

47 One practical way to distinguish content from context representations is to separate out mental
 48 (or neural) representations that drift rapidly (content; Manning et al., 2012; Polyn et al., 2005) or
 49 gradually (context; Folkerts et al., 2018; Howard et al., 2012; Lohnas et al., 2018; Long and Kahana,
 50 2018; Manning et al., 2011; Polyn et al., 2005). An intuition for why contextual representations
 51 might be expected to evolve gradually is laid out by Polyn and Kahana (2008). Essentially, slowly
 52 drifting thoughts may be used to “index” more rapidly drifting thoughts that are temporally
 53 proximal. This notion, that our brains maintain parallel mental representations drifting at different
 54 timescales, has been well-characterized in a series of elegant fMRI and ECoG studies by Uri
 55 Hasson and colleagues (Aly et al., 2018; Hasson et al., 2008; Honey et al., 2012; Lerner et al.,
 56 2011). Subsequent work by the same group has shown that this hierarchy of drifting mental
 57 representations plays a central role in how we remember continuous experiences and segment
 58 our experiences into discrete events (Baldassano et al., 2017). This body of work suggests that
 59 different brain regions reflect *temporal receptive windows* that characterize the timescale(s) to which
 60 each region is maximally sensitive. This research complements work on the *temporal window of*
 61 *integration* in the speech and perception literatures (e.g., Gauthier et al., 2012; van Wassenhove
 62 et al., 2007; Viemeister and Wakefield, 1991). The temporal window of integration refers to the
 63 duration over which physically distinct stimuli are treated (by some brain system or network)
 64 as a single percept or atomic unit. This systems-level work mirrors the discovery of *time cells*
 65 in the rodent hippocampus that respond preferentially to a specific time relative to a temporal
 66 reference (e.g., time elapsed since starting to run on an exercise wheel), whereby each cell appears
 67 to integrate incoming information at a different rate (MacDonald et al., 2011; Pastalkova et al.,
 68 2008) and the population activity serves to represent information at a range of timescales (Mau
 69 et al., 2018). Other work suggests that the lateral entorhinal cortex also plays a role in integrating
 70 information across a wide range of timescales (Tsao et al., 2018), supporting precise temporal
 71 recall (Montchal et al., 2019). In addition to this work suggesting that individual systems or
 72 neurons maintain representations at their preferred timescales, a number of studies also report
 73 neurons and small populations exhibiting *temporal multiplexing*, whereby information at different
 74 timescales is reflected in different features of an individual neuron’s (or population’s) activity

75 patterns, such as spike timing versus firing rate or local field potential oscillations at different
76 frequencies (Caruso et al., 2018; Knight and Eichenbaum, 2013; Panzeri et al., 2010; Watrous et al.,
77 2013).

78 The above ideas about parallel neural representations that reflect information at multiple
79 timescales have been formalized in a family of theoretical models over the past several decades (Di-
80 ana et al., 2007; Howard and Kahana, 2002; Howard et al., 2014; Polyn et al., 2009; Ranganath, 2018;
81 Sederberg et al., 2008; Shankar and Howard, 2010, 2012; Shankar et al., 2009). In turn, these models
82 were inspired in part by earlier theoretical work suggesting that temporally proximal experiences
83 become linked through their shared contextual features (Atkinson and Shiffrin, 1968; Estes, 1955).
84 Other recent work indicates that different aspects of the reinstatement process itself might also
85 occur at multiple timescales. For example, Linde-Domingo et al. (2019) found that high-level
86 semantic information was reinstated prior to low-level (e.g., perceptual) details in an associative
87 memory task. Collectively, these experimental and theoretical studies make two important con-
88 tributions to our understanding of how we remember our past. First, the studies suggest that
89 our ongoing experiences are “blurred out” in time by neural integrators (where different brain
90 structures or sub-structures integrate with different time constants). Second, the studies sug-
91 gest that representations that reflect different timescales become bound together such that when
92 we retrieve memories about our past, we reactivate representations that carry information at a
93 range of timescales. These reactivations explain how our brains reactivate prior contexts when we
94 remember the past.

95 Although reactivating context is often *described* as a sort of mental time travel back to the one
96 specific moment being remembered, there is some apparent inconsistency between this description
97 and how this process is actually characterized mathematically or measured neurally. In particular,
98 both the theory and neural measures described in the preceding paragraphs characterize episodic
99 recall (mathematically) as reflecting the reactivation of a weighted blend of thoughts from a *range*
100 of previously experienced times– not to any one *specific* time. This is not simply a matter of
101 imprecision in mental time travel (i.e., a reflection of the low effective resolution with which we
102 can revisit specific moments from our past). Rather, when we retrieve the thoughts associated

103 with a range of times, it is conceptually more like we are simultaneously visiting *many* of our prior
104 experiences, analogous to a quantum wave function spreading its probability mass over space
105 (see *Defining the quantum memory wave function metaphor and comparing it to the mental time*
106 *travel metaphor*). The above human and rodent studies suggest that the brain maintains parallel
107 representations of ongoing experience, each reflecting information at a different timescale. The
108 notion of thinking of any one particular moment from the past is incompatible with this multiple
109 timescales view, in that our thoughts and brain activity patterns reflect information from a range
110 of moments, even when we are not specifically engaged in remembering. If what we consider to
111 be a “moment” or “now” is guided by our neural representations of our ongoing experiences, then
112 this implies we are continually experiencing *many* “nows.” If our thoughts are continually spread
113 across a range of times, the notion of mentally revisiting any particular moment loses its meaning.

114 Where the mental time travel framing breaks down most notably is when we consider scenarios
115 in which the particular blend of timepoints reflected in one’s current thoughts do not come from
116 temporally contiguous events. Understanding how a new experience fits in with our broader
117 understanding might require integrating information gleaned from many experiences separated
118 in time, each with their own contextual attributes. For example, having an important conversation
119 with a friend might require bringing in information acquired during other interactions with that
120 friend, other conversations about related or similarly important topics, general knowledge about
121 communication styles, etc. This notion of bringing a range of prior experiences to bear on guiding
122 our ongoing behaviors is also related to a literature on *situation models* (for review see Ranganath
123 and Ritchey, 2012). That we can integrate information across experiences suggests that the inte-
124 gration of incoming information at different rates need not be a purely mechanical process, but
125 rather may reflect a deeper functional role (also see Brigard et al., 2018). When we simultaneously
126 reactivate thoughts related to separated moments from our past, this provides a means of associ-
127 ating and relating those temporally distinct experiences. This is how we can integrate information
128 across the *many* discrete experiences relevant to “now.”

129 Defining the *quantum memory wave function* metaphor and com- 130 paring it to the *mental time travel* metaphor

131 Unlike in classical physics, where matter dutifully obeys Newton’s laws, quantum physics postu-
132 lates a different (though potentially complementary) set of rules governing how matter behaves
133 and interacts. Perhaps the most famous equation in quantum physics is Schrödinger (1926)’s
134 equation describing the quantum *wave function* (this work was also inspired by de Broglie, 1924).
135 Schrödinger’s equation describes, in probabilistic terms, the location of a particle in space, at a
136 given time. In other words, it defines a function, $\Psi(x, t)$, such that the probability that a particle
137 is located at position x at time t is equal to $|\Psi(x, t)|^2$. The central idea is that, although it may
138 not be possible to know or predict the particle’s precise position at a given moment, the wave
139 function tells us how its location’s probability mass extends over space. In turn, this allows us
140 to use the wave function to predict behaviors or constraints of the system(s) in which the particle
141 is behaving and interacting. Within the domain of episodic memory, for the purposes of describ-
142 ing the *quantum memory wave function* metaphor, we can conceptualize the moments along our
143 *autobiographical timeline* (Arzy et al., 2009) as the spatial coordinates (x) in Schrödinger’s equation.
144 Then our mental state at time t may be conceptualized as spreading its probability mass over our
145 autobiographical timeline. In other words, like Schrödinger’s equation, many episodic memory
146 models (e.g., Howard and Kahana, 2002; Polyn et al., 2009; Shankar and Howard, 2012) define a
147 function describing how our mental state at time t reflects the blends of thoughts we experienced
148 at other moments of our lives. Under this metaphor, the effect of mental operations we perform
149 with respect to our autobiographical timeline, such as remembering past experiences or predicting
150 what will happen in the future, is to shift the probability mass of our quantum memory wave func-
151 tion (Fig. 1). This metaphor is also related to work on quantum probability decision models (e.g.,
152 Khrennikov et al., 2018) in that it extends similar concepts to the domain of episodic memory.

153 The *mental time travel* metaphor is defined by Tulving as “the type of awareness experienced
154 when one thinks back to a *specific moment* in one’s personal past and consciously recollects some
155 prior episode or state as it was previously experienced” (Fig. 1B; Wheeler et al., 1997, emphasis

Figure 1: **Memory wave functions.** **All panels.** Each panel displays a quantum memory wave function for a hypothetical individual at a particular moment (time $t = 0$ indicated by ▲). The x -axes correspond to the position (t) along the individual's autobiographical timeline. The y -axes display how much the individual's thoughts at time $t = 0$ include information about the individual's past (or future) perceived or predicted experiences (i.e., at other times). Different colors denote different amounts of temporal blurring. All curves have been normalized to have a maximum value of 1 and a minimum value of 0. **A. Contextual drift (no explicit reinstatement).** The blend of thoughts leading up to $t = 0$ reflects a recency-weighted average of information from the past. **B. Mental time travel.** This constitutes a special case of the quantum memory wave function framework, whereby the individual's thoughts reflect a recency-weighted blend of prior thoughts, along with thoughts specific to one remembered prior experience (occurring at the timepoint indicated by the ★ symbol). **C. Quantum memory waves.** In this example, the individual remembering the moment indicated by the ★ symbol brings back a blend of thoughts similar to those shown in Panel B. In addition, they bring to mind information pertaining to other conceptually related (but temporally non-contiguous) moments, indicated by ★ symbols. **D. Generating predictions about the future.** To predict what will happen during a series of future events (◇ symbols) the individual spreads their thoughts over their prior experiences according to the expected relevance of those experiences to the future events. Local peaks in the curve denote prior experiences with high expected relevance. The left part of the panel (prior to ▲) illustrates how the individual's thoughts are spread out over their past, and the right part of the panel illustrates how the individual's thoughts are spread out over their (predicted) future.

156 added). Whereas mental time travel refers to mentally revisiting one specific moment in addition
157 to now (potentially with some temporal blurring), the quantum memory wave framework implies
158 that episodic memory entails spreading our awareness over *many* moments (Fig. 1C). In this
159 way, the key distinction between these framings is whether remembering entails simultaneously
160 revisiting one single moment versus many moments. The two frameworks converge when we
161 engage in *chronesthesia* (Tulving, 2002a)– i.e., when we are consciously aware of subjective time.
162 Analogous to how measuring the physical location of a particle collapses its quantum wave function
163 at the moment of measurement, engaging in chronesthesia can collapse our quantum memory wave
164 function such that our thoughts temporarily converge on the moment we are revisiting. This wave
165 function expands again when our mental state changes in the next moment.

166 The quantum memory wave and mental time travel metaphors also diverge in explaining how
167 we might predict what will happen in the future. The quantum memory wave framing suggests
168 that our predictions about the future draw from information pertaining to a blend of our prior
169 experiences, weighted by their expected relevance to the future event in question (Fig. 1D). To the
170 extent that the mental time travel metaphor assumes that we can mentally visit just one other time
171 (or time range) in addition to now (e.g., Szpunar and Tulving, 2011), it is not clear how this sort
172 of integration over temporally separated prior experiences might occur. Specifically, integrating
173 over those experiences would seem to require some mechanism for associating and simultaneously
174 revisiting them while also casting some part of our thoughts forward to the future.

175 While the studies presented in the *Introduction and overview* (on temporal receptive windows, the
176 temporal window of integration, time cells, and temporal multiplexing) present strong empirical
177 evidence that our brains maintain parallel representations at multiple timescales, questions remain
178 about how or whether we experience those representations at a conscious level. For example, are
179 we consciously aware of spreading our mental states over many moments from our past? Or is our
180 subjective experience of remembering more like mentally time traveling back to a single “hybrid”
181 moment that blends information from several of our (actual) prior experiences?

182 The temporal dynamics of ongoing experience and how we re- 183 member it

184 Modern context-based theories of episodic memory posit that the current state of mental context
185 serves to determine which information from our past may be relevant to us and therefore reactiv-
186 ated (e.g., Polyn et al., 2009). Although these theories were largely developed and tested in the
187 domain of list-learning paradigms, conceptually similar principles might also underlie real-world
188 memory. For example, prior experiences with shared or overlapping contextual properties (e.g.,
189 other experiences that occurred in similar spatial or social settings, shared similar goals, etc.) could
190 be leveraged to form schemas or situation models that help guide behaviors and expectations
191 according to the current perceived context (Aly et al., 2018; Baldassano et al., 2018; Ranganath and
192 Ritchey, 2012).

193 Unlike random word lists, naturalistic experiences contain *event boundaries* whereby contextual
194 cues change sharply from one moment to the next, implying a change in the current situation (e.g.,
195 shifting from the quiet peace of working on a manuscript in the early morning to the frenzied
196 chaos of helping a newly awoken child get ready for school). A number of studies (using non-
197 random lists and more naturalistic stimuli) have found that the way we remember past experiences
198 can be shaped by these event boundaries. For example, controlling for elapsed time, memory is
199 impaired for information that occurred prior to a change in the current event or situation (e.g.,
200 Ezzyat and Davachi, 2011; Manning et al., 2016; Radvansky and Copeland, 2006; Swallow et al.,
201 2011, 2009). The temporal dynamics of contextual change also underlie how we judge the amount
202 of time elapsed since a prior reference point (Block and Reed, 1978; Sahakyan and Smith, 2014).
203 These findings suggest that the rapid contextual or situational changes that define these event
204 boundaries shape how we experience and remember (DuBrow and Davachi, 2016), similarly to how
205 spatial boundaries (e.g., environmental barriers; Brunec et al., 2018; McKenzie and Buzsáki, 2016)
206 shape how we experience and remember spatial environments and layouts, or how conceptual
207 boundaries (e.g., distinctions between semantic categories; Brunec et al., 2018) shape how we
208 experience and remember conceptual information.

209 Event boundaries are one example of the broader temporal covariance structure that is charac-
210 teristic of ongoing experience. Whereas classic approaches to studying memory (e.g., list-learning
211 paradigms and other trial-based experiments) encourage researchers to treat each moment, trial,
212 or stimulus as separable from the rest of experience, context-based theories of memory (including
213 situation models and event-based models) posit that each moment derives meaning through its
214 deep ties to other related (but potentially temporally separated) experiences. This rich tapestry of
215 interacting moments and experiences forms a scaffolding for interpreting our new experiences and
216 for retrieving information concerning our past. According to this view, the reactivation of this rich
217 tapestry of contextual details enfolding the remembered event (i.e., how the event relates to the
218 rest of one's experiences across timescales) is even more central to our subjective experience than
219 the specific details of what occurred during that event! This begs the question: which thoughts,
220 from which times, do we reactivate when we remember our past experiences?

221 **How is our past replayed when we remember?**

222 Our ongoing experiences can cue us to remember information about our past. This can happen
223 explicitly (e.g., during a conversation about a specific prior event) or implicitly (e.g., when some-
224 thing happening now reminds you of what happened before). Context-based theories of memory
225 posit that the particular memories that our ongoing experiences cue will be related (contextually,
226 semantically, functionally, etc.; e.g., Polyn et al., 2009). One potential challenge to studying how
227 we remember the past relates to detangling which links between our experiences are specifically
228 driven by our memory systems versus which are imposed "externally" through the similarity
229 structure of those experiences. For example, consider your daily commute into work. If you
230 were to verbally recount a specific day's commute later, one might expect some aspects of your
231 description to apply to other days' commutes as well (e.g., other commutes where you took a
232 similar route or travelled via similar modes of transportation, listened to similar music along the
233 way, travelled at a similar time, etc.). But to what extent could one hope to separate out descriptive
234 matches driven solely by similarities in those commuting experiences, versus true reactivations (in
235 memory) of specific memories of those other prior commutes?

236 One insight into this distinction comes from a recent study by Heusser et al. (2018a), which
 237 analyzed data collected by Chen et al. (2017) who had participants watch an episode of the BBC
 238 television show *Sherlock* and then verbally recount what happened in the episode. The authors
 239 used topic models (Blei et al., 2003) applied to human-generated annotations of each camera shot
 240 in the episode to characterize the episode’s content with a timeseries of high-dimensional semantic
 241 feature vectors (see *Defining the geometries of ongoing thoughts and memories*). Each shot’s annotation
 242 contained the following information: a brief narrative description of what was happening, the
 243 location where the scene took place, whether that location was indoors or outdoors, the names
 244 of all on-screen characters, the names of the character(s) who were in focus in the camera shot,
 245 the name(s) of any character(s) who spoke during the scene, the camera angle of the shot, a
 246 transcription of any on-screen text, and whether or not there was any music present in the auditory
 247 background (full details may be found in Chen et al., 2017). The resulting *topic vectors* describe
 248 the mix of “themes” present in each moment of the episode (Fig. 5B). When characterized in
 249 this way, the episode’s content exhibits periods of relative stability, followed by moments of
 250 rapid change. This interplay between stability and change is visible as “blocks” in the temporal
 251 correlation matrix (Fig. 2A); Heusser et al. (2018a) performed detailed analyses of these putative
 252 *events* and how they are remembered. Notably, the off-diagonal entries of the temporal correlation
 253 matrix (which reflect content similarity between temporally distant events in the episode) are
 254 nearly all close to 0. This is surprising, because the episode itself comprises a murder mystery with
 255 recurring motifs that must be associated in the viewer’s mind in order to make sense of the episode.
 256 Nonetheless, when considering solely the semantic content of each moment of the episode, those
 257 long timescale associations are nearly entirely absent. To the extent that this finding extends to
 258 real-world experiences, it comports with Heraclitus’ notion that one cannot step in the same river
 259 twice (Heraclitus, BCE).

260 Although the specific mix of contents in each moment of the episode is nearly unique, partic-
 261 ipants’ subjective experiences of the episode entail weaving together those temporally separated
 262 moments into a cohesive narrative. This is reflected in how people verbally recount the episode
 263 later. If we examine the temporal autocorrelation matrices of participants’ recalls of the television

Figure 2: **Temporal correlation matrices of the dynamic content of a television episode and people’s recalls of that episode.** The panels of this figure are adapted from Heusser et al. (2018a), where full details may be found. **A. Temporal correlations exhibited by the content of the television episode.** Each moment of an episode of the BBC television show *Sherlock* has been characterized using a topic model (Blei et al., 2003) fit to detailed human-generated annotations between each scene cut. The correlations between the topic vectors of a given pair of timepoints reflects the similarity in thematic content at those timepoints. **B. Temporal correlations exhibited by the (average) content of participants’ recalls of the episode.** The topic model used to characterize the episode in Panel A has been applied to each sentence from each participants’ recall transcript. The temporal correlation matrices describing each participants’ recalls were re-sampled using cubic interpolation to have 100 rows and columns prior to averaging the resampled matrices across participants. **C. Temporal correlations in the content of individual participants’ recalls of the episode.** Each sub-panel displays the temporal correlation matrix for one experimental participant.

episode (characterized using the same topic model used to capture the content of the episode; Figs. 2B, C), we see a different structure than that of the episode itself. Although individual participants' recalls also exhibit periods of content stability punctuated by rapid change (visible in the blocky structure of the matrices in Fig. 2C), the off-diagonal correlations between the thematic content of temporally distant events are substantially stronger in participants' recalls than in the original episode. Another way of characterizing this phenomenon is to look at the autocorrelation functions (describing the correlation between different moments' topic vectors as a function of their temporal distance) for the episode versus participants' recalls (Fig. 3A). Whereas the autocorrelations in the original episode's content rapidly fall to 0, autocorrelations in the content of participants' recalls appear to asymptote around an average of 0.2. These analyses indicate that the way participants *recount* temporally separated events imposes additional similarity on those events, beyond what was present in the original events themselves. One can then ask: is this simply a matter of imprecision in recall, either with respect to participants' recalls themselves, or with respect to the quality with which they are characterized via semantic feature vectors? Or might there be a functional role of that additional imposed similarity structure?

To understand which specific events are described (during recall) as more similar than what their original experiences (as characterized by their topic vectors) would suggest, it can be helpful to examine individual recalls of particular events throughout the episode. Figure 3B displays, for a representative example recalled event, which other parts of the episode were characterized in a similar way by the participants when they recounted the episode later (black and colored lines), and which other parts of the episode were described in a similar way in moment-by-moment (non-recall) human annotations (gray line). In the example event explored in the figure, Sherlock Holmes and John Watson (main characters) are bantering, and then are interrupted when Sherlock realizes that another person has committed suicide. Later the viewer learns that the suicides were actually murders. The language used by human annotators describing the moment-by-moment contents of the event used a similar mix of themes only for temporally nearby parts of the episode (gray line). However, the way participants described what happened when they *recounted* the episode later revealed a much different pattern. Specifically, their language mirrored that used

292 when they were recalling other related murders, and other moments when Sherlock and John were
293 bantering with or about each other (Fig. 3C).

294 The way they described the example target event did *not* mirror the way they described other
295 events that were not directly narratively linked with the target event, even when those other
296 events overlapped semantically. For example, thematic elements of events involving unrelated
297 violence, or Sherlock and John interacting with other characters about other subjects, tended not
298 to feature in participants' recountings of the target event. Taken together, this suggests that these
299 specific events were associated in participants' memories, perhaps contributing to (or reflecting)
300 their understanding that those events were linked in the episode's narrative.

301 One concrete example of this phenomenon is that early in the episode, including during the
302 target event, Sherlock and John appear to be investigating a series of suicides. Later in the episode,
303 it becomes clear that the apparent suicides are actually murders. The annotations refer to early
304 deaths (including that shown in Fig. 3B) as suicides, whereas participants often refer to those early
305 deaths as "murders" even though they could not have known that they were murders early on in
306 the episode. In this way, participants' experiences *subsequent to* a target remembered event may
307 distort its memory by incorporating information learned during other experiences.

308 A recall density function analogous to the one displayed in Figure 3B could, in principle, be
309 characteristic of memory for *any* past event or experience. For example, the peaks of the func-
310 tion could provide insight into which events, from which moments along one's autobiographical
311 timeline, are functionally associated in memory. Nonetheless, a potential challenge to testing this
312 hypothesis in the laboratory is that in many memory studies there is no particular reason for
313 participants to functionally associate the stimuli. For example, there is often no "deeper narrative"
314 to make sense of in list-learning studies or in most standard trial-based memory experiments.
315 However, random word list learning experiments might happen to contain (for a given list) se-
316 mantically related words. Indeed, a number of studies indicate that participants spontaneously
317 form associations between semantically related words on random word lists, even when those
318 related words are separated in time (e.g., Manning and Kahana, 2012; Manning et al., 2012; Wixted
319 and Rohrer, 1994). This suggests that our memory systems pick up on statistical regularities, even

Figure 3: Content drift in a television episode and participants' recalls of that episode. **A. Autocorrelations in the episode's content and in the content of participants' recalls.** The black line denotes the average correlations between the topic vectors from pairs of moments in the episode as a function of their temporal separation (times are relative to the full length of the episode). Each colored line denotes an analogous autocorrelation function, but for each individual participant's recalls of the episode (times are relative to the full length of each participant's recalls). The error ribbons denote 95% confidence intervals, taken across all moments in the episode (or in participants' recalls of the episode). **B. Recall density functions for one recalled event.** Each curve reflects the correlations between the topic vector at each (relative) moment of the episode (or participants' recalls of the episode), and the topic vector of an event beginning approximately $\frac{1}{3}$ through the episode (relative time: 0.328). The gray curve denotes these correlations for the topic vectors derived from annotations of the episode itself. The black line denotes the average correlations across all participants' recalls. Local maxima are marked with colored dots. Local maxima closer than 2.5% of the total recall duration to a higher local maximum were excluded (this constraint served to eliminate very nearby peaks that corresponded to duplicate events), as were maxima with peaks with correlation coefficients lower than 0.1 (this threshold was chosen arbitrarily; this constraint served to eliminate events that were only very weakly associated with the example event). The colored lines denote the correlations for each individual participant's recalls. **C. Local maxima in the average recall density function.** Representative screen captures from the scenes at each local maxima in Panel B are displayed (dot colors match those in Panel B). The text in each screen capture displays a brief summary of what happened in each scene. The colored outline denotes the reference scene being recalled.

320 in ostensibly “random” stimulus sequences, and that we leverage those regularities when we recall
321 our past (Polyn et al., 2009). Other work suggests that our memory systems may also leverage these
322 regularities to parse our ongoing experiences into discrete events (Schapiro et al., 2013; Shapiro,
323 2019). Are these examples in the literature on random list learning studies, indicating that our
324 memory systems leverage “accidental” structure in random sequences, a reflection of the same
325 fundamental principle driving the memory reactivation function displayed in Figure 3B?

326 It is difficult to explicitly test whether (apparent) semantic reinstatement in word list learning,
327 as reflected in people’s tendencies to semantically cluster recalls of random word lists, reflects
328 integration of information across the experiences of studying those words. However, remembering
329 word lists does not specifically *require* integrating information across words in order to gain a deep
330 understanding of the list. By design, there *is* no deeper meaning of random word lists. One way
331 to characterize the meaning reflected by how an experience unfolds in time is through models
332 that describe the geometric path that the experience takes through an appropriate representational
333 space.

334 **The shapes of remembered experiences**

335 The study by Heusser et al. (2018a) described above suggests one way of characterizing the dynamic
336 content of complex naturalistic stimuli like television episodes. The sequence of topic vectors over
337 the course of the episode traces out a *topic trajectory* through a high-dimensional feature space.
338 The sequence of topic vectors for a given participant’s recalls traces out an analogous trajectory.
339 This characterization provides a geometric framework for comparing the *shapes* that reflect how
340 naturalistic experiences unfold over time, to the shapes of how we remember those experiences
341 later (Fig. 4). One can then use the geometric transformations needed to map an experience’s
342 trajectory onto the trajectory of its later remembering to characterize the quality and content of
343 memory.

344 In addition to serving as a framework for quantifying the dynamic content of experiences, these
345 trajectories may be visualized by projecting word embedding representations of the events onto

Figure 4: **The shape of an unfolding experience.** **A. The trajectory through topic space taken by a television episode.** The topic trajectory taken by a 50 minute television episode, projected onto two dimensions. Each dot indicates an event identified using a Hidden Markov Model; the dot colors denote the order of the events (early events are in red; later events are in blue), and the connecting lines indicate the transitions between successive events. **B. The trajectory through topic space taken by participants’ recalls of the television episode.** The average two-dimensional trajectory captured by 17 participants’ recalls of the episode, with the same format and coloring as the trajectory in Panel A. The arrows reflect the directions of the average transitions taken by participants whose trajectories crossed that part of topic space; blue arrows denote directions of reliable agreement across participants. Figure adapted from Heusser et al. (2018a), where additional details may be found.

lower dimensional spaces (e.g., Heusser et al., 2018b); for additional discussion and counterpoints see Jolly and Chang (2018). Heusser et al. (2018a) applied this approach to the same dataset reflected in Figures 2 and 3, revealing another key property of how we remember. Although there are substantial individual differences in the specific words people use to describe their shared experience of watching the same television episode, and how much detail they recollected, the overall shapes of those recalled experience trajectories are highly similar across people. This has potentially exciting implications for what it means to communicate memories to other people: to effectively communicate, we often need to convey the “essence” (gross overall shape) of our experience. It also has implications for how we learn by analogy: perhaps events that share schemas (e.g., Baldassano et al., 2018) unfold in similar ways (e.g., have similar shapes). Learning which shapes go with which schemas could provide a scaffolding for rapid learning and generalization of new experiences, to the extent that those new experiences match a previously learned schema (shape).

In their paper, Heusser et al. (2018a) use the geometric shapes of experience trajectories (through word embedding space) to functionally distinguish high-level essential details of the experience

360 from low-level (non-essential) details. When the trajectory exhibits a sharp change in direction,
361 this indicates a rapid shift in the thematic content of the episode. Heusser et al. (2018a) used
362 hidden Markov models to detect these rapid shifts and segment the trajectories into discrete
363 *events*. The high-level (essential) details of the episode are reflected in the shape defined by the
364 sequence of coordinates of its events. These are captured by low spatial frequency properties of
365 the trajectory’s shape. Low-level details are characterized by (typically smaller-scale) within-event
366 geometric patterns, which are captured by the high spatial frequency properties of the trajectory’s
367 shape. A key finding of that work was that people who watched the episode recalled high-level
368 details reliably, and in a similar way across people, whereas people varied substantially in how
369 they remembered low-level details. This suggests that our memory systems may prioritize certain
370 types of information (e.g., high-level details) about our experiences over others (e.g., low-level
371 details) as we encode our experiences into memories or when we recount our prior experiences.

372 ***The matching problem***

373 The trajectories through word embedding space displayed in Figure 4 demonstrate that topic
374 models are sufficiently expressive and robust to provide one solution to the *matching problem* in
375 memory analysis. The matching problem refers to discovering an alignment in the labels between
376 experienced and remembered events. Traditional memory paradigms most commonly studied
377 over the past century (Kahana, 2012) ask participants to study lists of *items* (often randomly selected
378 words) and then engage in a memory task, such as verbally recalling any words the participant
379 remembers studying, in any order. Solving the matching problem in such experiments is trivial:
380 if we treat each studied word on a list as an “event,” then we can assume that the participant is
381 remembering a word’s event when they verbalize that word during their memory test. In other
382 words, matching entails finding the exact match between a word the participant is verbalizing and
383 that same word on the list that they studied.

384 For real-world memory phenomena, and in many naturalistic memory paradigms, solving the
385 matching problem is more complex. People often experience real-world events (and scenes from
386 movies or stories) as rich sensory experiences that derive deep meaning from surrounding mo-

ments, other prior experiences, emotions, goals, high-order understanding of the current situation, and so on. Many of these aspects of our real-world experiences are not immediately verbalizable—and even when we are able to describe some specific aspects of what we are experiencing, there are typically many possible ways we could have done so (e.g., varying in level of detail, specific wording, objectivity versus subjectivity, contextual details, etc.). Further, because the way we experience each moment depends in part on our unique prior experiences, even ostensibly shared experiences can *seem* different to each individual. The way we recount a given experience is often a reflection of our own idiosyncratic experience of the original event, and is colored and distorted in an idiosyncratic way by our memory systems. This distorted memory must then be verbalized or communicated, which is itself an idiosyncratic (and potentially lossy, in the information theoretic sense) process. Solving the matching problem under these constraints poses a substantial theoretical and methodological challenge. Specifically, matching requires building explicit models of the dynamic conceptual content of our experiences and memories, and then estimating when and how that experienced and remembered content aligns. The approach featured in Figure 4 uses a word embedding space that is defined by a topic model (Blei et al., 2003). Each coordinate in this space reflects the conceptual content underlying a set of words, such that two different text passages that describe the same conceptual content (but using different words, levels of detail, etc.) are mapped onto nearby points in the word embedding space. When these models are sufficiently expressive and robust to solve the matching problem, by modeling our experiences and memories within the same word embedding space, we can see a clear geometric alignment between the trajectories that our experiences (Fig. 4A) and memories (Fig. 4B) take through word embedding space (also see Heusser et al., 2018a).

Defining the geometries of ongoing thoughts and memories

Defining geometric spaces that can reliably capture the conceptual contents of our thoughts and experiences has become a major focus of computational science. One family of models that has shown particular promise over the past several decades is broadly referred to as *word embedding models* (Fig. 5). The fundamental goal of these models is to take as input a sequence of words (text)

414 and produce as output one or more feature vectors whose coordinates reflect their conceptual
415 (semantic) content. For example, when two conceptually similar documents are passed through
416 these models, the goal is to produce two feature vectors that are closer (e.g., in Euclidean distance)
417 than the feature vectors for two dissimilar documents.

418 Early word embedding models such as *Latent Semantic Analysis* (LSA; Deerwester et al., 1990;
419 Landauer and Dumais, 1997) derived text embeddings using word co-occurrence statistics. The logic
420 driving this approach posits that documents that contain similar mixtures of words are likely to be
421 conceptually similar, and words that appear in similar documents are also likely to be conceptually
422 similar. To quantify these intuitions, Deerwester et al. (1990) first constructed *word count matrices*
423 (Fig. 5A) for several text corpora (each comprising many thousands of documents). Each entry
424 in these matrices denotes the number of times a given document (row) contained a given word
425 (column). They then used the Singular Value Decomposition (SVD) to decompose the number-of-
426 documents $\times$ number-of-words word count matrix into the product of a number-of-documents
427 $\times K$ *document embeddings matrix*, a $K \times K$ *singular values matrix*, and a $K \times$ number-of-unique-words
428 *word embeddings matrix*. The rows of the document embeddings matrix reflect the similarities
429 between documents, and the columns of the word embeddings matrix reflect similarities between
430 words.

431 A related approach, termed *Word Association Spaces* (WAS; Nelson et al., 2004; Steyvers et al.,
432 2004) replaces the word count matrix with data from a large free-association experiment. In the
433 experiment, each participant was asked to free-associate (i.e., say the first word that came to mind)
434 in response to a series of cue words. The authors then constructed a number-of-cue-words $\times$
435 number-of-unique-response-words matrix whose entries reflected the proportion of times each
436 cue word (row) produced a given response word (column) across participants. After using SVD
437 to decompose this *word frequency matrix*, they obtained low-dimensional embeddings of each cue
438 word and response word. In one sense, WAS might provide a more accurate and behaviorally
439 relevant reflection of how people think about the meanings of words than LSA, since nearby WAS
440 vectors reflect either cue words that produced similar responses, or response words that were
441 evoked by similar cues. In other words, WAS derives embeddings directly from people’s semantic

judgement behaviors. However, because WAS is based on a large amount of experimental data, updating the embeddings requires collecting new data, which can be slow and costly. As the target application or population diverges from the experimental population (e.g., over time, across cultures, demographics, etc.) the WAS embeddings may become less accurate or relevant. By contrast, LSA may be re-fit relatively quickly to any new corpus, enabling the resulting embeddings to stay current. Further, because LSA may be applied to any corpus, word embeddings may be easily computed and/or compared across languages, times, subject areas, etc.

Topic models, such as *Latent Dirichlet Allocation* (LDA; Blei et al., 2003) represent an attempt to characterize the latent *topics* or themes that documents comprise (Fig. 5B). Whereas matrix factorization approaches represent semantic meanings using low-dimensional approximations of word co-occurrences or free association frequencies, topic models leverage these properties of documents (or behaviors) to explain *why* words co-occur. LDA defines a topic as a distribution of weights over words in the vocabulary. For example, if one takes the *vocabulary* to be the top 10,000 most commonly used words in English (excluding stop words that are devoid of explicit semantic meaning such as “and,” “or,” “the,” “but,” etc.), then each topic comprises a vector of 10,000 non-negative weights that sum to one. As shown in the cartoon on the left of the Panel (where the four top-weighted words from each of five hypothetical topics are displayed), a fruit-related topic (shown in purple) might weight heavily on words like “apple,” “banana,” “pear,” and “orange,” whereas words like “cat,” “dog,” “dolphin,” and “lobster” might receive relatively small weights. Given a set of K topic vectors, topic models compute the blend of topics most probably reflected in a given document—i.e., its *topic proportions*. These topic proportion estimates are made by examining how different topics weight each of the words in the document, and by assuming that documents tend to reflect few topics (rather than many topics). Given only a word count matrix as input (just as in LSA), topic models automatically discover a $K \times \text{number-of-words}$ *topics matrix*, which describes how each topic (row) weights each word in the vocabulary (column), and a $\text{number-of-documents} \times K$ *topic proportions matrix*, which describes each document’s topic proportions. The rows of the topic proportions matrix may be treated as low-dimensional embeddings of the set of documents.

A Matrix factorization models

B Topic models

C Deep learning models

Figure 5: Three approaches to text embedding. **A. Matrix factorization models.** Per-document (or per-stimulus) word counts and/or frequencies are collected into a word count matrix (left). To compute K -dimensional embeddings of the documents and words, the word count matrix is factorized (typically by computing its Singular Value Decomposition) into the product of a number-of-documents $\times K$ *document embeddings matrix*, a $K \times K$ *singular values matrix*, and a $K \times$ number-of-unique-words *word embeddings matrix*. **B. Topic models.** Fitting a topic model to a word count matrix (e.g., Panel A) reveals a set of latent *topics*, or themes that pervade each document. The topics are specified in a $K \times$ number-of-words *topics matrix*, and the per-document themes are specified in a number-of-documents $\times K$ *topic proportions matrix*. A given document's topic proportions reflect the specific mix of themes that are reflected in that document (bar graph on the left). These proportions are discovered automatically by estimating the topic from which each word in the document is most probably drawn (highlighted words in *Document text*). **C. Deep learning models.** A wide variety of deep learning approaches to text embedding have been proposed. Broadly, these approaches often combine a common set of elements, including *transformers* (diagram adapted from Viswani et al., 2017), *deep averaging networks* (diagram adapted from Iyyer et al., 2015), *recurrent neural networks* (diagram adapted from Iyyer et al., 2015), and *deep autoencoders* (diagram adapted from Yousefi-Azar and Hamey, 2017).

470 A key advantage of topic models over matrix factorization models is their ability to explain
471 potentially ambiguous word usage. For example, consider the use of the word “bat” in the
472 following illustrative passage:

473 *“As the slavering mutant vampire **bat** swooped low to attack, its eyes ablaze, the explorer grunted*
474 *and took a mighty swing of her trusty aluminum baseball **bat**. The powerful strike connected*
475 *with an echoing metallic clink, stunning the wretched creature before it scrambled about on the*
476 *cave floor and flew away. She didn’t even **bat** an eye— as the Dark Wizard Chiroptera continued*
477 *to draw strength from the mystic crystals, such things had become increasingly common of late,*
478 *nearly to the point of mundanity...”*

479 At different points in the passage, “bat” refers to an animal, a piece of sports equipment, and an eye-
480 related verb. Because each of these uses of the word appears in the same (hypothetical) document,
481 models driven solely by word co-occurrence statistics would blend these different meanings such
482 that the embedding of “bat” was somewhat similar to other animals *and* other pieces of sports
483 equipment *and* other eye-related verbs. By contrast, treating each document as reflecting a blend
484 of topics allows these uses to be disambiguated. If other documents that mention the word “bat”
485 describe vampires, swooping, or other mammals, but *not* sports equipment or eye-related verbs,
486 then a topic model could explain this pattern by inferring that multiple topics place substantial
487 weight on the word “bat.” Another strength of topic models is their interpretability. Because the
488 models directly output each document’s topic proportions along with each topic’s weights on all
489 words in the vocabulary, the general themes that each topic reflects are often intuitive and the
490 per-document topic proportions may be directly related to those themes.

491 Both the matrix factorization models and topic model described above are each examples of
492 *bag of words* models. Given a document, these models will generate the same embeddings for any
493 random re-ordering of the same words. This over-simplifies the intended conceptual content of
494 documents, where word order often conveys important meaning. An alternative framing of the
495 text embedding problem is to use *sequence* models that explicitly account for word order effects.
496 Most modern sequence-based models are implemented using deep neural networks (Fig. 5C).

497 The *Neural Probabilistic Language Model* (NPLM; Bengio et al., 2003) was one of the earliest deep
498 learning-based sequence models. NPLM derives word embeddings for each word in the vocab-
499 ulary (such that nearby embeddings reflect semantically related words), and then models word
500 sequences according to their meanings. This model was inspired in part by n -gram models (Brown
501 et al., 1992), which estimate the probability of the next word in a sequence given the n words
502 that came before it. Effectively, NPLM attempts to learn arbitrarily long sequences. In a related
503 approach, Mnih and Hinton (2009) proposed the *Scalable Hierarchical Distributed Language Model*,
504 which learns sequences of variable length n -grams to achieve greater explanatory power.

505 The next major advances in word embeddings came from attempts to more precisely capture the
506 geometric relations between words. For example, if each word’s semantic meaning is captured by a
507 feature vector, then in principle one could perform vector algebra on different words’ embeddings
508 to produce new concepts that combine their meanings. For example, subtracting the vector for
509 “man” from that of “king” and then adding the vector for “woman” might result in a new feature
510 vector similar to that of “queen” (Mikolov et al., 2013b). Mikolov et al. (2013a) developed the deep
511 learning model *word2vec* to explicitly optimize the accuracy of these vector operations. Whereas
512 *word2vec* is a bag of words model, *Global Vectors for Word Representation* (GloVe; Pennington et al.,
513 2014) extends this approach to model each word’s local (temporal) context (also see Peters et al.,
514 2018).

515 Several recent models have begun to move beyond single word-level embeddings. For example,
516 *SkipThought* (Kiros et al., 2015), *Universal Sentence Encoder* (Cer et al., 2018), *InferSent* (Conneau
517 et al., 2018), and *RandSent* (Wieting and Kiela, 2019) each compute sentence-level embeddings. The
518 central insight of these approaches is that capturing the nuanced meaning of a particular concept
519 often requires *several* words that have been organized into a complete sentence or sentence-like
520 grammatical construct.

521 Cutting-edge approaches to word embeddings that have achieved the best-to-date performance
522 on modern natural language processing benchmarks are built using enormous deep neural net-
523 works often with many billions of parameters. Because these models are very computationally
524 expensive to fit, an important discovery has been that many layers in these models may be pre-

525 trained without substantially affecting performance, even on new tasks. One of the first of this new
526 class of model, *Bidirectional Encoder Representations from Transformers* (BERT; Devlin et al., 2018),
527 uses a pre-trained deep neural network, whose weights in a single output layer may be fine-tuned
528 to suit a wide range of new tasks. The latest additions to this class of models are called *Generative*
529 *Pre-trained Transformers* (GPT; Brown et al., 2020; Radford et al., 2018, 2019).

530 **Extensions to spatial, semantic, affective, and social representations**

531 The way we navigate through *thought space* (i.e., word embedding spaces; Figs. 4, 5) differs from how
532 we navigate in physical spatial environments. First, we cannot teleport in real space, whereas our
533 thoughts can exhibit rapid jumps between different non-contiguous regions of word embedding
534 space. Second, when two people visit the same fixed physical landmark, they necessarily visit
535 the same physical location. When two people describe a shared experience, their idiosyncratic
536 thoughts, goals, prior experiences, etc. can lead them to process and remember the experience
537 differently, thereby visiting different locations within thought space. Third, each time we revisit
538 a fixed physical landmark, we (by definition) revisit the same spatial location. In other words, it
539 is possible to visit the same spatial location multiple times. However, we cannot revisit our prior
540 experiences in this way. Rather, we only get to experience each moment once; subsequent “visits”
541 to a prior experience must occur through a different medium than the original experience (e.g., our
542 memory, another person’s memory, a physical recording, a written account, etc.). Fourth, we can
543 revisit different aspects of the same prior experience on different “visits” – e.g., when we remember
544 a specific event we can focus on the physical occurrence in one remembering, the emotional
545 occurrence in another remembering, the social implications in another remembering, etc. We can
546 also revisit the same experience in different levels of detail or depth. There are no deep analogs of
547 these phenomena in spatial navigation. Finally, *physically* (at a macro scale), we are always located
548 at a single moment and in a single location. The main argument of the quantum memory wave
549 function framework I propose here is that we may be *mentally* spread over many times or locations
550 simultaneously. For example, this provides a way for us to integrate information from multiple
551 prior experiences into a single decision or conceptualization, even if those experiences were not

552 temporally contiguous.

553 Many of these differences between physical and mental navigation also apply to other mental
554 domains, including how we *conceptualize* space (e.g., when we trace out a potential route in
555 our mind; Arzy and Schacter, 2019; Gauthier et al., 2018; Gauthier and van Wassenhove, 2016).
556 Just as we mentally spread ourselves across a distribution of times when we remember, we can
557 perform similar mental operations when we conceptualize and retrieve other types of information.
558 Further, these distributions may be estimated using existing methods. For example, consider
559 how we use pattern classifiers (Norman et al., 2006) to connect neural activity patterns to specific
560 semantic (Manning et al., 2012; Mitchell et al., 2008; Polyn et al., 2005), spatial (Miller et al., 2013),
561 affective (Chang et al., 2018), and social (Meyer et al., 2018) content being mentally manipulated.
562 A prevailing approach is to label specific brain activity patterns as each reflecting *one particular*
563 concept or mental state. Variability in neural activity over repeated experiences or rememberings
564 can result in classifiers spreading the probability masses of their predictions over multiple possible
565 mental states, including non-contiguous mental states. Traditional approaches then “clean up”
566 this distribution over possible mental states by taking its average (expected), maximum, or modal
567 value. However, following the “quantum wave function” logic described in this manuscript,
568 these multi-peaked response patterns may provide a mechanism for associating non-contiguous
569 representations. What are some implications of this view for non-memory domains?

570 One example of how the full distribution over possible mental states can provide insights into
571 cognition may be seen in how the brain encodes spatial locations. Theoretical (e.g., Gerstner and
572 Abbott, 1997) and empirical (e.g., Pfeiffer and Foster, 2013) work has demonstrated how place
573 cell receptive fields may be modulated according to environmental features such as obstacles and
574 the locations of rewarded goals. In other words, although place cells encode an animal’s location
575 within an environment, even for a fixed cell and environment there is some variability with respect
576 to where precisely within that environment the cell modulates its firing rate. Further, a number
577 of studies have suggested that some neurons respond to distinct non-contiguous places (Derdik-
578 man et al., 2009; Fenton et al., 2008; Grieves et al., 2020; Lee et al., 2019; Rich et al., 2014) and
579 times (Pastalkova et al., 2008). Taken together, these studies suggest that variability in the set of

580 locations and/or times to which a neuron responds is not merely a reflection of noise or imprecision.
581 Rather, this variability might have some functional relevance in terms of facilitating the learning
582 of paths around obstacles and towards goals.

583 The notion that variability and “uncertainty” can be informative also has implications for how
584 our brains might efficiently represent information. Recent work suggests that adversarial networks
585 trained to learn a generative model of natural images (e.g., Gatys et al., 2016; Isola et al., 2017)
586 may be used to refine estimates of decoded visual stimuli (St-Yves and Naselaris, 2018). Related
587 manifold learning approaches that “snap” decoded estimates to the nearest surface location on a
588 learned manifold may be used in a similar way to decode sequences of movie frames (Heusser
589 et al., 2018b). Taken together, these findings suggest that not all *possible* images (and by extension,
590 all possible stimuli) are equally likely. Rather, the set of stimuli likely to be encountered live
591 within subspaces or on manifolds in the relevant representational space. The geometries of these
592 manifolds provide insights into the statistical structure of the stimuli. They also provide potential
593 insights into how the brain represents those stimuli, for example to the extent that the geometries of
594 these manifolds may be leveraged to improve brain decoding. The shapes of these manifolds can
595 potentially provide insights into how our thoughts are distributed over representational spaces.

596 **Concluding remarks**

597 When we remember an experience, we bring back thoughts related to a range of other conceptually
598 and functionally related experiences. This can serve an important functional role: it allows us to
599 integrate information across temporally separated experiences. Although this view is captured
600 mathematically by context-based theories of episodic memory, it is *not* specifically captured by the
601 conceptual notion of mental time travel that is often used to describe episodic recall. The alter-
602 native conceptualization presented here, that when we remember our past we reinstate thoughts
603 associated with a weighted blend of our prior experiences, also has links with non-temporal or
604 episodic domains such as semantic, spatial, affective, and social representations. Characterizing
605 how our thoughts reflect *distributions* over representational spaces is central to understanding how

606 we think about ourselves, our experiences, other people, and the world around us.

607 **Acknowledgements**

608 I thank Luke Chang, Janice Chen, Paxton Fitzpatrick, Andrew Heusser, Christopher Honey, Caro-
609 line Lee, Talia Manning, Lucy Owen, Matthijs van der Meer, and Kirsten Ziman for feedback and
610 scientific discussions. I also thank Janice Chen, Yuan Chang Leong, Christopher Honey, Chung
611 Yong, Kenneth Norman, and Uri Hasson for sharing the data re-analyzed here (Chen et al., 2017),
612 and Paxton Fitzpatrick and Andrew Heusser for contributing analysis code that I adapted from a
613 prior publication (Heusser et al., 2018a). This work was supported in part by NSF EPSCoR Award
614 Number 1632738. The content is solely my responsibility and does not necessarily represent the
615 official views of my supporting organizations.

616 **Data and code availability**

617 All analysis code and data used in the present manuscript may be found [here](#).

618 **References**

- 619 Aly, M., Chen, J., Turk-Browne, N. B., and Hasson, U. (2018). Learning naturalistic temporal
620 structure in the posterior medial network. *Journal of Cognitive Neuroscience*, 30(9):1345–1365.
- 621 Arzy, S., Adi-Japha, E., and Blanke, O. (2009). The mental time line: an analogue of the mental
622 number line in the mapping of life events. *Consciousness and Cognition*, 18:781–785.
- 623 Arzy, S. and Schacter, D. L. (2019). Self-agency and self-ownership in cognitive mapping. *Trends in*
624 *Cognitive Sciences*, 23(6):476–486.
- 625 Atkinson, R. C. and Shiffrin, R. M. (1968). Human memory: A proposed system and its control

processes. In Spence, K. W. and Spence, J. T., editors, *The psychology of learning and motivation*,
volume 2, pages 89–105. Academic Press, New York.

Baldassano, C., Chen, J., Zadbood, A., Pillow, J. W., Hasson, U., and Norman, K. A. (2017).
Discovering event structure in continuous narrative perception and memory. *Neuron*, 95(3):709–
721.

Baldassano, C., Hasson, U., and Norman, K. A. (2018). Representation of real-world event schemas
during narrative perception. *Journal of Neuroscience*, 38(45):9689–9699.

Bengio, Y., Ducharme, R., and Vincent, P. (2003). A neural probabilistic language model. *Journal of*
Machine Learning Research, 3:1137–1155.

Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of Machine*
Learning Research, 3:993 – 1022.

Block, R. A. and Reed, M. A. (1978). Remembered duration: evidence for a contextual-change
hypothesis. *Journal of Experimental Psychology: Human Learning and Memory*, 4:656–665.

Brigard, F. D., Hanna, E., Jacques, P. L. S., and Schacter, D. L. (2018). How thinking
about what could have been affects how we feel about what was. *Cognition and Emotion*,
doi:10.1080/02699931.2018.1478280.

Brown, P. F., deSouza, P. V., and Mercer, R. L. (1992). Class-based n -gram models of natural
language. *Computational Linguistics*, 18(4):467–480.

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam,
P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R.,
Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S.,
Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. (2020).
Language models are few-shot learners. *arXiv*, 2005.14165.

Brunec, I. K., Moscovitch, M. M., and Barense, M. D. (2018). Boundaries shape cognitive represen-
tations of spaces and events. *Trends in Cognitive Sciences*, 22(7):637–650.

651 Caruso, V. C., Mohl, J. T., Glynn, C., Lee, J., Willett, S. M., Zaman, A., Ebihara, A. F., Estrada, R.,
652 Freiwald, W. A., Tokdar, S. T., and Groh, J. M. (2018). Single neurons may encode simultaneous
653 stimuli by switching between activity patterns. *Nature Communications*, 9(1):2715.

654 Cer, D., Yang, Y., Kong, S. Y., Hua, N., Limtiaco, N., John, R. S., Constant, N., Guajardo-Cespedes,
655 M., Yuan, S., Tar, C., Sung, Y.-H., Strope, B., and Kurzweil, R. (2018). Universal sentence encoder.
656 *arXiv*, 1803.11175.

657 Chang, L. J., Jolly, E., Cheong, J. H., Rapuano, K., Greenstein, N., Chen, P.-H., and Manning,
658 J. R. (2018). Endogenous variation in ventromedial prefrontal cortex state dynamics during
659 naturalistic viewing reflects affective experience. *bioRxiv*, doi.org/10.1101/487892.

660 Chen, J., Leong, Y. C., Honey, C. J., Yong, C. H., Norman, K. A., and Hasson, U. (2017). Shared
661 memories reveal shared structure in neural activity across individuals. *Nature Neuroscience*,
662 20(1):115.

663 Conneau, A., Kiela, D., Schwenk, H., Barrault, L., and Bordes, A. (2018). Supervised learning of
664 universal sentence representations from natural language inference data. *arXiv*, 1705.02364.

665 de Broglie, L. (1924). *Recherches sur la théorie des quanta*. PhD thesis, Migration-université en cours
666 d’affectation.

667 Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., and Harshman, R. (1990). Indexing
668 by latent semantic analysis. *Journal of the American Society for Information Science*, 41(6):391–407.

669 Derdikman, D., Whitlock, J., Tsao, A., Fyhn, M., Hafting, T., Moser, M., and Moser, E. (2009).
670 Fragmentation of grid cell maps in a multicompartment environment. *Nature Neuroscience*,
671 12(10):1325–1332.

672 Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). Bert: pre-training of deep bidirectional
673 transformers for language understanding. *arXiv*, 1810.04805.

- 674 Diana, R. A., Yonelinas, A. P., and Ranganath, C. (2007). Imaging recollection and famil-
675 iarity in the medial temporal lobe: a three-component model. *Trends in Cognitive Sciences*,
676 doi:10.1016/j.tics.2007.08.001.
- 677 DuBrow, S. and Davachi, L. (2016). Temporal binding within and across events. *Neurobiology of*
678 *Learning and Memory*, 134:107–114.
- 679 Estes, W. K. (1955). Statistical theory of spontaneous recovery and regression. *Psychological Review*,
680 62:145–154.
- 681 Ezzyat, Y. and Davachi, L. (2011). What constitutes an episode in episodic memory? *Psychological*
682 *Science*, 22(2):243–252.
- 683 Fenton, A. A., Kao, H.-Y., Neymotin, S. A., Olypher, A., Vayntrub, Y., Lytton, W. W., and Ludvig, N.
684 (2008). Unmasking the CA1 ensemble place code by exposures to small and large environments:
685 more place cells and multiple, irregularly arranged, and expanded place fields in the larger
686 space. *The Journal of Neuroscience*, 28(44):11250–11262.
- 687 Folkerts, S., Rutishauser, U., and Howard, M. W. (2018). Human episodic mem-
688 ory retrieval is accompanied by a neural contiguity effect. *Journal of Neuroscience*,
689 doi.org/10.1523/JNEUROSCI.2312-17.2018.
- 690 Gatys, L. A., Ecker, A. S., and Bethge, M. (2016). Image style transfer using convolutional neural
691 networks. *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2414–2423.
- 692 Gauthier, B., Eger, E., Hesselmann, G., Giraud, A.-L., and Kleinschmidt, A. (2012). Temporal tuning
693 properties along the human ventral visual stream. *The Journal of Neuroscience*, 32(41):14433–14441.
- 694 Gauthier, B., Pestke, K., and van Wassenhove, V. (2018). Building the arrow of time...over time: a se-
695 quence of brain activity mapping imagined events in time and space. *Cerebral Cortex*, 29(10):4398–
696 4414.
- 697 Gauthier, B. and van Wassenhove, V. (2016). Cognitive mapping in mental time travel and mental
698 space navigation. *Cognition*, 154:55–68.

699 Gerstner, W. and Abbott, L. F. (1997). Learning navigational maps through potentiation and
700 modulation of hippocampal place cells. *Journal of Computational Neuroscience*, 4:79 – 94.

701 Grieves, R. M., Jeddi-Ayoub, S., Mishchanchuk, K., Liu, A., Renaudineau, S., and Jeffery, K. J.
702 (2020). The place-cell representation of volumetric space in rats. *Nature Communications*, 11(789).

703 Hassabis, D., Kumaran, D., Vann, S. D., and Maguire, E. A. (2007). Patients with hippocampal
704 amnesia cannot imagine new experiences. *Proceedings of the National Academy of Science USA*,
705 107:1726–1731.

706 Hasson, U., Yang, E., Vallines, I., Heeger, D. J., and Rubin, N. (2008). A hierarchy of temporal
707 receptive windows in human cortex. *Journal of Neuroscience*, 28(10):2539–2550.

708 Heraclitus (535 BCE - 475 BCE). *You cannot step twice into the same river; for other waters are continually*
709 *flowing in*. Fragment 41; Quoted by Plato in *Cratylus*.

710 Heusser, A. C., Fitzpatrick, P. C., and Manning, J. R. (2018a). How is experience transformed into
711 memory? *bioRxiv*, 409987.

712 Heusser, A. C., Ziman, K., Owen, L. L. W., and Manning, J. R. (2018b). HyperTools: a Python
713 toolbox for gaining geometric insights into high-dimensional data. *Journal of Machine Learning*
714 *Research*, 18(152):1–6.

715 Honey, C. J., Thesen, T., Donner, T. H., Silbert, L. J., Carlson, C. E., Devinsky, O., Doyle, J. C.,
716 Rubin, N., Heeger, D. J., and Hasson, U. (2012). Slow cortical dynamics and the accumulation of
717 information over long timescales. *Neuron*, 76:423–434.

718 Howard, M. W. and Kahana, M. J. (2002). A distributed representation of temporal context. *Journal*
719 *of Mathematical Psychology*, 46:269–299.

720 Howard, M. W., MacDonald, C. J., Tiganj, Z., Shankar, K. H., Du, Q., Hasselmo, M. E., and H., E.
721 (2014). A unified mathematical framework for coding time, space, and sequences in the medial
722 temporal lobe. *Journal of Neuroscience*, 34(13):4692–4707.

- Howard, M. W., Viskontas, I. V., Shankar, K. H., and Fried, I. (2012). Ensembles of human MTL neurons “jump back in time” in response to a repeated stimulus. *Hippocampus*, 22:1833–1847.
- Isola, P., Zhu, J.-Y., Zhou, T., and Efros, A. A. (2017). Image-to-image translation with conditional adversarial networks. *IEEE Conference on Computer Vision and Pattern Recognition*, pages 5967–5976.
- Iyyer, M., Manjunatha, V., Boyd-Graber, J., and Daumé III, H. (2015). Deep unordered composition rivals syntactic methods for text classification. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*.
- Jolly, E. and Chang, L. J. (2018). The flatland fallacy: moving beyond low dimensional thinking. *Topics in Cognitive Science*, doi:10.1111/tops.12404.
- Kahana, M. J. (2012). *Foundations of Human Memory*. Oxford University Press, New York, NY.
- Khrennikov, A., Basieva, I., Pothos, E. M., and Yamato, I. (2018). Quantum probability in decision making from quantum information representation of neuronal states. *Scientific Reports*, 8(16225).
- Kiros, R., Zhu, Y., Salakhutdinov, R., Zemel, R., Torralba, A., Urtasun, R., and Fidler, S. (2015). Skip-thought vectors. *Advances in Neural Information Processing Systems*, pages 3294–3302.
- Knight, R. T. and Eichenbaum, H. (2013). Multiplexed memories: a view from human cortex. *Nature Neuroscience*, 16(3):257–258.
- Landauer, T. K. and Dumais, S. T. (1997). A solution to Plato’s problem: the latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104:211–240.
- Lee, J. S., Briguglio, J., Romani, S., and Lee, A. K. (2019). The statistical structure of the hippocampal code for space as a function of time, context, and value. *bioRxiv*, 10.1101/615203.
- Lerner, Y., Honey, C. J., Silbert, L. J., and Hasson, U. (2011). Topographic mapping of a hierarchy of temporal receptive windows using a narrated story. *Journal of Neuroscience*, 31(8):2906–2915.

- 748 Linde-Domingo, J., Treder, M. S., Kerrén, C., and Wimber, M. (2019). Evidence that neural informa-
749 tion flow is reversed between object perception and object reconstruction from memory. *Nature*
750 *Communications*, 10(179):1–13.
- 751 Lohnas, L. J., Duncan, K., Doyle, W. K., Thesen, T., Devinsky, O., and Davachi, L. (2018). Time-
752 resolved neural reinstatement and pattern separation during memory decisions in human hip-
753 pocampus. *Proceedings of the National Academy of Science USA*, doi.org/10.1073/pnas.1717088115.
- 754 Long, N. M. and Kahana, M. J. (2018). Hippocampal contributions to serial-order memory. *Hip-*
755 *pocampus*, page <https://doi.org/10.1002/hipo.23025>.
- 756 MacDonald, C., Lepage, K., Eden, U., and Eichenbaum, H. (2011). Hippocampal “time cells” bridge
757 the gap in memory for discontinuous events. *Neuron*, 71(4):737–749.
- 758 Manning, J. R., Hulbert, J. C., Williams, J., Piloto, L., Sahakyan, L., and Norman, K. A. (2016).
759 A neural signature of contextually mediated intentional forgetting. *Psychonomic Bulletin and*
760 *Review*, 23(5):1534–1542.
- 761 Manning, J. R. and Kahana, M. J. (2012). Interpreting semantic clustering effects in free recall.
762 *Memory*, 20(5):511–517.
- 763 Manning, J. R., Norman, K. A., and Kahana, M. J. (2015). The role of context in episodic memory.
764 In Gazzaniga, M., editor, *The Cognitive Neurosciences, Fifth edition*, pages 557–566. MIT Press.
- 765 Manning, J. R., Polyn, S. M., Baltuch, G., Litt, B., and Kahana, M. J. (2011). Oscillatory patterns
766 in temporal lobe reveal context reinstatement during memory search. *Proceedings of the National*
767 *Academy of Sciences, USA*, 108(31):12893–12897.
- 768 Manning, J. R., Sperling, M. R., Sharan, A., Rosenberg, E. A., and Kahana, M. J. (2012). Spon-
769 taneously reactivated patterns in frontal and temporal lobe predict semantic clustering during
770 memory search. *The Journal of Neuroscience*, 32(26):8871–8878.
- 771 Mau, W., Sullivan, D. W., Kinsky, N. R., Hasselmo, M. E., Howard, M. W., and Eichenbaum, H.

772 (2018). The same hippocampal CA1 population simultaneously codes temporal information
 773 over multiple timescales. *Current Biology*, 28(1):1499–1508.

774 McKenzie, S. and Buzsáki, G. (2016). *Micro-, Meso- and Macro-Dynamics of the Brain, Research and*
 775 *Perspectives in Neurosciences*, chapter Hippocampal mechanisms for the segmentation of space
 776 by goals and boundaries. Springer.

777 Meyer, M. L., Davachi, L., Ochsner, K. N., and Lieberman, M. D. (2018). Evidence that default
 778 network connectivity during rest consolidates social information. *Cerebral Cortex*, 10.1093/cer-
 779 cor/bhy071.

780 Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013a). Efficient estimation of word representations
 781 in vector space. *arXiv*, 1301.3781.

782 Mikolov, T., Yih, W.-T., and Zweig, G. (2013b). Linguistic regularities in continuous space word
 783 representations. *Proceedings of the NAACL-HLT*, pages 746–751.

784 Miller, J. F., Neufang, M., Solway, A., Brandt, A., Trippel, M., Mader, I., Hefft, S., Merkow, M., Polyn,
 785 S. M., Jacobs, J., Kahana, M. J., and Schulze-Bonhage, A. (2013). Neural activity in human hip-
 786 pocampal formation reveals the spatial context of retrieved memories. *Science*, 342(6162):1111–
 787 1114.

788 Mitchell, T., Shinkareva, S., Carlson, A., Chang, K., Malave, V., Mason, R., and Just, M. (2008).
 789 Predicting human brain activity associated with the meanings of nouns. *Science*, 320(5880):1191.

790 Mnih, A. and Hinton, G. (2009). A scalable hierarchical distributed language model. *Advances in*
 791 *Neural Information Processing Systems*.

792 Montchal, M. E., Reagh, Z. M., and Yassa, M. A. (2019). Precise temporal memories are supported by
 793 the lateral entorhinal cortex in humans. *Nature Neuroscience*, doi.org/10.1038/s41593-018-0303-1.

794 Nelson, D. L., McEvoy, C. L., and Schreiber, T. A. (2004). The University of South Florida free asso-
 795 ciation, rhyme, and word fragment norms. *Behavior Research Methods, Instruments and Computers*,
 796 36(3):402–407.

797 Norman, K. A., Polyn, S. M., Detre, G. J., and Haxby, J. V. (2006). Beyond mind-reading: multi-voxel
798 pattern analysis of fMRI data. *Trends in Cognitive Sciences*, 10(9):424–430.

799 Panzeri, S., Brunel, N., Logothetis, N. K., and Kayser, C. (2010). Sensory neural codes using
800 multiplexed temporal scales. *Trends in Neurosciences*, 33(3):111–120.

801 Pastalkova, E., Itskov, V., Amarasingham, A., and Buzsáki, G. (2008). Internally generated cell
802 assembly sequences in the rat hippocampus. *Science*, 321:1322 – 1327.

803 Pennington, J., Socher, R., and Manning, C. D. (2014). GloVe: global vectors for word representation.
804 In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*.

805 Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., and Zettlemoyer, L. (2018).
806 Deep contextualized word representations. *arXiv*, 1802.05365.

807 Pfeiffer, B. E. and Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to
808 remembered goals. *Nature*, 497:74–79.

809 Polyn, S. M. and Kahana, M. J. (2008). Memory search and the neural representation of context.
810 *Trends in Cognitive Sciences*, 12:24–30.

811 Polyn, S. M., Natu, V. S., Cohen, J. D., and Norman, K. A. (2005). Category-specific cortical activity
812 precedes retrieval during memory search. *Science*, 310:1963–1966.

813 Polyn, S. M., Norman, K. A., and Kahana, M. J. (2009). A context maintenance and retrieval model
814 of organizational processes in free recall. *Psychological Review*, 116(1):129–156.

815 Radford, A., Narasimhan, K., Salimans, T., and Sutskever, I. (2018). Improving language under-
816 standing by generative pre-training. Unpublished whitepaper.

817 Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. (2019). Language models are
818 unsupervised multitask learners. *OpenAI Blog*, 1(8).

819 Radvansky, G. A. and Copeland, D. E. (2006). Walking through doorways causes forgetting:
820 situation models and experienced space. *Memory & Cognition*, 34(5):1150 – 1156.

- 821 Ranganath, C. (2018). Time, memory, and the legacy of Howard Eichenbaum. *PsyArXiv*,
822 10.31234/osf.io/9tcrp.
- 823 Ranganath, C. and Ritchey, M. (2012). Two cortical systems for memory-guided behavior. *Nature*
824 *Reviews Neuroscience*, 13:713 – 726.
- 825 Rich, P. D., Liaw, H.-P., and Lee, A. K. (2014). Large environments reveal the statistical structure
826 governing hippocampal representations. *Science*, 345(6198):804–817.
- 827 Sahakyan, L. and Smith, J. R. (2014). A long time ago, in a context far, far away: Retrospective
828 time estimates and internal context change. *Journal Experimental Psychology: Learning, Memory*
829 *and Cogntion*, 40(1):86–93.
- 830 Schacter, D. L., Addis, D. R., and Buckner, R. L. (2007). Remembering the past to imagine the
831 future: the prospective brain. *Nature Reviews Neuroscience*, 8:657–661.
- 832 Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B., and Botvinick, M. M. (2013).
833 Neural representations of events arise from temporal community structure. *Nature Neuroscience*,
834 16:486–492.
- 835 Schrödinger, E. (1926). An undulatory theory of the mechanics of atoms and molecules. *The Physical*
836 *Review*, 28(6):1049–1070.
- 837 Sederberg, P. B., Howard, M. W., and Kahana, M. J. (2008). A context-based theory of recency and
838 contiguity in free recall. *Psychological Review*, 115(4):893–912.
- 839 Shankar, K. H. and Howard, M. W. (2010). Timing using temporal context. *Brain Research*, 1365(3-
840 17).
- 841 Shankar, K. H. and Howard, M. W. (2012). A scale-invariant internal representation of time. *Neural*
842 *Computation*, 24:134–193.
- 843 Shankar, K. H., Jagadisan, U. K. K., and Howard, M. W. (2009). Sequential learning using temporal
844 context. *Journal of Mathematical Psychology*, 53:474 – 485.

- 845 Shapiro, M. L. (2019). Time is just a memory. *Nature Neuroscience*, doi.org/10.1038/s41593-018-0331-
846 X.
- 847 St-Yves, G. and Naselaris, T. (2018). Generative Adversarial Networks conditioned on brain activity
848 reconstruct seen images. *bioRxiv*, <http://dx.doi.org/10.1101/304774>.
- 849 Steyvers, M., Shiffrin, R. M., and Nelson, D. L. (2004). Word association spaces for predicting
850 semantic similarity effects in episodic memory. In Healy, A. F., editor, *Cognitive Psychology and
851 its Applications: Festschrift in Honor of Lyle Bourne, Walter Kintsch, and Thomas Landauer*. American
852 Psychological Association, Washington, DC.
- 853 Swallow, K. M., Barch, D. M., Head, D., Maley, C. J., Holder, D., and Zacks, J. M. (2011). Changes in
854 events alter how people remember recent information. *Journal of Cognitive Neuroscience*, 23(5):1052
855 – 1064.
- 856 Swallow, K. M., Zacks, J. M., and Abrams, R. A. (2009). Event boundaries in perception affect
857 memory encoding and updating. *Journal of Experimental Psychology: General*, 138(2):236 – 257.
- 858 Szpunar, K. K. and Tulving, E. (2011). Varieties of future experience. In Bar, M., editor, *Predictions
859 in the brain: using our past to generate a future*. Oxford University Press.
- 860 Tsao, A., Sugar, J., Lu, L., Wang, C., Knierim, J. J., Moser, M.-B., and Moser, E. I. (2018). Integrating
861 time from experience in the lateral entorhinal cortex. *Nature*, 561:57–62.
- 862 Tulving, E. (1972). Episodic and semantic memory. In Tulving, E. and Donaldson, W., editors,
863 *Organization of Memory.*, pages 381–403. Academic Press, New York.
- 864 Tulving, E. (1983). *Elements of Episodic Memory*. Oxford, New York.
- 865 Tulving, E. (2002a). Chronesthesia: conscious awareness of subjective time. In Stuss, D. T. and
866 Knight, R. T., editors, *Principles of frontal lobe function*. Oxford University Press.
- 867 Tulving, E. (2002b). Episodic memory: from mind to brain. *Annual Review of Psychology*, 53:1–25.

- 868 van Wassenhove, V., Grant, K. W., and Poeppel, D. (2007). Temporal window of integration in
869 auditory-visual speech perception. *Neuropsychologia*, 45:598–607.
- 870 Viemeister, N. F. and Wakefield, G. H. (1991). Temporal integration and multiple looks. *Journal of*
871 *the Acoustic Society of America*, 90(2):858–865.
- 872 Viswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., and Kaiser, L. (2017).
873 Attention is all you need. In *31st Conference on Neural Information Processing Systems*.
- 874 Watrous, A. J., Tandon, N., Conner, C. R., Pieters, T., and Ekstrom, A. D. (2013). Frequency-specific
875 network connectivity increases underlie accurate spatiotemporal memory retrieval. *Nature Neu-*
876 *roscience*, 16(3):349–356.
- 877 Wheeler, M. A., Stuss, D. T., and Tulving, E. (1997). Toward a theory of episodic memory: the
878 frontal lobes and autonoetic consciousness. *Psychological Bulletin*, 121(3):331–354.
- 879 Wieting, J. and Kiela, D. (2019). No training required: exploring random encoders for sentence
880 classification. *arXiv*, 1901.10444.
- 881 Wixted, J. T. and Rohrer, D. (1994). Analyzing the dynamics of free recall: An integrative review
882 of the empirical literature. *Psychonomic Bulletin and Review*, 1:89–106.
- 883 Yousefi-Azar, M. and Hamey, L. (2017). Text summarization using unsupervised deep learning.
884 *Expert Systems with Applications*, 68:93–105.