Journal of Memory and Language 102 (2018) 130–141

Contents lists available at ScienceDirect

Journal of Memory and Language

journal homepage: www.elsevier.com/locate/jml

Patterns of forgetting

Jerry S. Fisher*, Gabriel A. Radvansky

University of Notre Dame, United States

ARTICLE INFO

Keywords:
Retention
Forgetting
Event models
Word lists
Narratives

ABSTRACT

In memory research, forgetting is largely assumed to occur following a relatively consistent forgetting curve. However, recent work in our lab suggests that there is a shift in the pattern of forgetting after a retention interval of about seven days. Moreover, work on narrative comprehension has shown that information at different levels of representation show different patterns of forgetting. Much of the existing work on patterns of forgetting (a) does not allow one to assess changes in forgetting because much of the data is collected either prior to or after seven days, but not sufficiently bridging this period of time, and (b) does not consider patterns of forgetting for different levels of memory representation. In this study, memory for a list of words and narrative texts was assessed up to 12 weeks after initial learning. We observed that memory for the word list showed some forgetting early on, followed by an abrupt loss after about seven days. Moreover, for the narrative text, surface form memories were forgotten to around chance level after about an hour, whereas textbase level memories were retained until about seven days when memory suddenly dropped to around chance levels, much like the word list memories. In contrast to this, memory at the event model level remained high throughout, although there was some forgetting over time. To account for this pattern of retention and forgetting, a simulation was developed as a proof of concept to illustrate our theoretical interpretation.

Introduction

One of the most enduring findings in the study of human memory is Ebbinghaus's (1885) *retention curve* (often called the *forgetting curve*). This curve is a negatively accelerating function in which the majority of forgetting occurs soon after learning, with less information being forgotten as time progresses. The basic assumption of most memory researchers is that this is a continuous function, which progresses in a relatively constant pattern in long-term memory. The aim of the current study is to more finely assess the forgetting of information over long periods of time. This was done for both (a) a standard set of memoranda in memory research, namely word lists, as well as for (b) more complex sets of information, namely memory for narrative texts.

We first review some recent evidence from our lab suggesting a shift in memory retention and forgetting after about seven days. After this, we discuss memory for different levels of representation of narrative texts, as well as two studies that have looked at memory for text at these levels at longer retention intervals. Next, we present the current study. After this we, present a simulation as a proof of concept for the patterns of retention that we observe.

Forgetting curve shifts

There has been a fair amount work on the retention curve over the years. The bulk of this work has been focused on determining the nature of this function, such as whether it is an exponential or a power function, with the consensus being that it is somewhat better described by a power function (Wixted & Ebbesen, 1991). Some have explored the idea that a power function description is a result of an averaging of memory performance across many trials or observations (Anderson & Tweney, 1997; Averell & Heathcote, 2011; but see Wixted & Ebbesen, 1997). If the decay of individual memory traces were exponential, averaging across multiple exponential functions is best fit by a power function (Murre & Chessa, 2011), namely $M = at^b$, where M is memory, t is time, a is a constant, and b is the exponent, conveying the rate of forgetting over log time. The most important component for us here is the exponent, b. For this paper, we treat a retention curve as described by a power function as the default pattern and thus our null hypothesis for retention. Any deviation from this would need to be taken into account.

Although a power function is the default assumption, a recent analysis of existing retention data by Pettijohn and Radvansky (2017) and Csik and Radvansky (2018) suggests that there may be changes in the rate of forgetting over time, as defined by the exponent of the power function. For a power function, the exponent fit to a forgetting curve

* Corresponding author at: Department of Psychology, University of Notre Dame, Notre Dame, IN 46556, United States.
E-mail address: jerry.s.fisher.111@nd.edu (J.S. Fisher).

https://doi.org/10.1016/j.jml.2018.05.008
Received 26 July 2017; Received in revised form 22 May 2018

should remain constant throughout the process of retention and forgetting (Wixted, 2004). Thus, if one fits a power function to early time points, the expectation is that future time points would be predicted to follow that same rate. Pettijohn and Radvansky tested this prediction by using data from 23 published studies, including 45 experiments, from Ebbinghaus (1885) to the present, which had five or more retention intervals. Power functions were fit to the data from the first four retention intervals to predict performance for later retention intervals. They found that while retention intervals that were seven days or shorter showed better memory than predicted, those intervals that were greater than seven days showed the opposite pattern, with faster than predicted forgetting. This suggests that the rate of forgetting, as captured by the power function exponent, may increase after seven days.

Moreover, the study by Csik and Radvansky (2018) explicitly addressed the size of the exponent of the power function, which can provide an index of the rate of forgetting. Using data from 44 published studies, including 135 experiments, which had three or more retention intervals, the data were fit to a power function, and the exponent from that function was recorded. These exponents were then analyzed as a function of the longest retention interval for a given experiment. What was found was that from about one minute to one day, there is a decrease in the rate of forgetting, with exponents approaching zero. This parallels the idea that LTP in the hippocampus takes several hours to complete and may be aided by sleep. From about one day to about nine days, the rate of forgetting remains largely stable, with exponents being similarly closer to zero. Then, after this time, for exponents of functions lasting up to years later, the rate of forgetting increases with the values largely moving further away from zero.

Although this idea that there is a transition somewhere around seven days is seen when analyzing across retention studies, there are few studies that have sufficient time points both prior to and after this transition period of seven days within their retention range. Thus, this current study directly bridges this gap for different material types.

Levels of representation

Another factor that can affect memory retention is the type of representation that is involved. Here we make use of narrative memories because they can be readily divided into three levels of representation. Specifically, these are the surface form, the textbase, and the event model levels (Van Dijk, Kintsch, & Van Dijk, 1983). The surface form is memory for the verbatim words and syntax that were used. This type of memory is typically very short lasting, often just a few minutes (Sachs, 1967). The textbase is memory for the propositional idea units that are conveyed in a text independent of the wording. Thus, a verbatim sentence from a text and a paraphrase that conveys the same meaning would map onto the same textbase representation. Finally, while the surface form and textbase representations are memories for the text itself, the event model is a representation of the situation described by the text (Glenberg, Meyer, & Lindem, 1987). This referential representation contains information that was conveyed by the text, as well as any inferences a reader may draw based on their prior knowledge.

Schmalhofer and Glavanov (1986) developed a method of assessing each of these levels of representation in memory using a signal detection analysis. For each critical sentence from a text, there are four types of recognition probe sentences: (a) the *verbatim* sentence that actually appeared in the text, (b) a *paraphrase* of a verbatim sentence, which did not actually appear in the text, although that idea was conveyed, (c) an *inference* sentence that conveys an idea that was likely generated by readers using their world knowledge, and (d) a *wrong* sentence that, while generally thematically consistent with the text, is inconsistent with the events described by the text.

The measure of the surface form compares memory for verbatim probes (hits) with memory for paraphrases (false alarms). Both probe types capture idea units that were presented in the text, but only the verbatim convey the actual words and syntax used. The measure of the textbase compares memory for the paraphrases (hits) with memory for the inferences (false alarms). Both of these probe types were never actually mentioned in the text, and are consistent with the situation described by the text, but only the paraphrase conveys idea units that were actually present in the text. Finally, the measure of the event model compares memory for the inference probes (hits) with memory for the wrong probes (false alarms). Neither of these probes types were mentioned in the text, nor did they convey idea units actually present in the text; they were both thematically consistent with the text, however, but only the inferences were actually consistent with the situation described by the text. This method of assessing different levels of text representation has been used widely over the years (Bohay, Blakely, Tamplin, & Radvansky, 2011; Fletcher & Chrysler, 1990; Kintsch, Welsch, Schmalhofer, & Zimny, 1990; Narvaez, Radvansky, Lynchard, & Copeland, 2011; Radvansky, Copeland, & von Hippel, 2010; Radvansky, Copeland, & Zwaan, 2003; Radvansky, Gibson, & McNerney, 2014; Radvansky, Zwaan, Curiel, & Copeland, 2001; Zwaan, 1994).

Of particular interest here are two studies that have reported changes in the different levels of representation over different periods of time. The first is a study by Kintsch et al. (1990). Of the two experiments reported by Kintsch et al., we focus on the first. This experiment assessed memory for the surface form, textbase, and event model levels for people of four different retention groups: immediate testing after reading, and testing after 40 min, two days, or four days. What was found, as can be seen in Fig. 1, was that there was forgetting for the surface form and textbase measures, but there was no clear evidence of forgetting for the event model level. Given this surprising finding, it is important to (a) assess whether there is any forgetting at the event model level, (b) understand the pattern of forgetting, if there is any, and (c) provided an account for what gives rise to such a long-lasting representation of memory.

The second study of interest is by Radvansky et al. (2001). Of the two experiments they reported, we focus on the second, which assessed memory after a delay. This study compared younger and older adults on narrative memory at these three levels either immediately or after

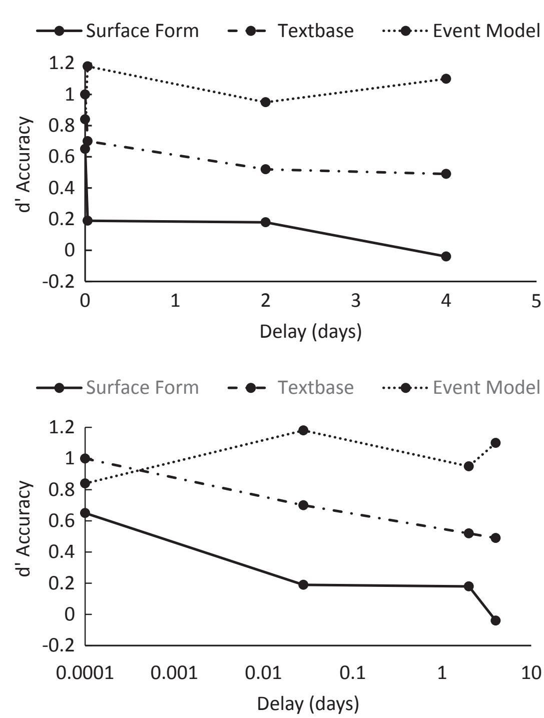

Fig. 1. Accuracy data from Kintsch et al. (1990) plotted in terms of a linear time scale (top) and a logarithmic time scale (bottom).