Title

A memory for goals model of prospective memory

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

Authors

Trafton, Greg
Harrison, Anthony

Publication Date

2023

Peer reviewed

A memory for goals model of prospective memory

J. Gregory Trafton (greg.trafton@nrl.navy.mil)
Naval Research Laboratory
Washington, DC 20375

Anthony M. Harrison (tony.harrison@nrl.navy.mil)
Naval Research Laboratory
Washington, DC 20375

Abstract

We present a novel model of prospective memory and fit it to data from a classic experimental paradigm (Einstein & McDaniel, 1990). Our model uses memory for goals (Altmann & Trafton, 2002) and elaboration with spreading activation to show how prospective intentions can be cued by perceptual cues. Our model also suggests how to resolve some of the controversies concerning prospective memory and aging.

Keywords: prospective memory; act-r; cognitive modeling; mind-wandering

Introduction

The vast majority of memory research has focused on retrospective memory, or memory of the past (Anderson, Bothell, Lebiere, & Matessa, 1998; Baddeley, 2012; Henson, 1998). However, memory for future intentions, or prospective memory is also extremely important, both from a theoretical and applied perspective and is much less studied. More formally, prospective memory is setting a goal or task to do in the future based on some cue or event (visual, aural, temporal). The classic prospective memory problem is to pick up a loaf of bread on the way home from work.

Ellis (1996) suggested there are five stages to successfully completing a prospective memory task.

Encoding An intent to perform a future action in a particular target context is created.

Retention There is a delay between the current time and the target context. The intention may be rehearsed (or not).

Initiation The agent perceives the target context, and a window of opportunity arises. The agent must recognize it as the target of an intention.

Execution The agent performs the stored action.

Completion The agent must modify its memory such that the next perception of the target does not lead to action.

Most experimental and real-life instances of prospective memory also have another step: The agent resumes the task it was working on before the prospective memory task. While this step is not strictly part of prospective memory, it provides some constraints and possible memory interactions within theories of prospective memory. Task resumption is also part of virtually every experiment on prospective memory.

Prospective memory is particularly difficult to model (Li & Laird, 2013). Li and Laird (2013) suggest that there is a circular knowledge dependency in cognitive modeling. The circular knowledge dependency problem is that if the memory target is in working memory, a memory retrieval is not needed, but if the memory target is not in working memory, the procedure to retrieve the target never applies. Because memory retrieval relies on procedural knowledge (Laird, Lebiere, & Rosenbloom, 2017), it becomes a bit of a circular problem.

Several other researchers have built computational models of prospective memory. Elio (2006) used a ACT-R's top-down spreading activation model and focused on modeling reaction time, not accuracy. Lebiere and Lee (2002) also used ACT-R's spreading activation to model the reaction time of the intention superiority effect. Li and Laird (2014) used Soar's (Laird, 2019) spreading activation mechanism to perform spontaneous retrievals to allow a computational agent to perform prospective memory tasks.

There are several interesting points to make on these models of prospective memory. First, all these models of prospective memory seem to rely on some form of spreading activation. Second, the models that focus on high cognitive fidelity (Elio, 2006; Lebiere & Lee, 2002) do not model the accuracy of prospective memory tasks, making them more difficult to apply to agents or robots (Trafton et al., 2013). In contrast, the model that focuses on functionality (Li & Laird, 2014) does not model human-level performance, but can perform prospective memory tasks. We also note that both empirical and computational accounts suggest that prospective memory has a strong episodic emphasis – people seem to remember prospective memory events in terms of discrete episodes (Li & Laird, 2014).

Our goal in this paper is to introduce a model that has high cognitive plausibility and that can perform prospective memory tasks by using episodic memory. Not surprisingly, it will use spreading activation as one of its primary mechanisms. Our model will be demonstrated with respect to one of the very first experimental demonstrations of prospective memory (Einstein & McDaniel, 1990).

Experiment: Einstein and McDaniel (1990)

Einstein and McDaniel (1990) were interested in providing an experimental demonstration of prospective memory. They were also interested in showing whether prospective mem-

In M. Goldwater, F. K. Anggoro, B. K. Hayes, & D. C. Ong (Eds.), *Proceedings of the 45th Annual Conference of the Cognitive Science Society*.

ory decreased as people aged, similar to how retrospective memory seems to decrease as people age (Rhodes, Greene, & Naveh-Benjamin, 2019; Luo & Craik, 2008).

Method

We present a brief description of Einstein and McDaniel's experiment, but a complete description can be found in Einstein and McDaniel (1990), experiment 2.

Design

The experiment was a 2x2 between design with age (young, elderly) and the familiarity of the prospective memory cue (familiar, unfamiliar).

Participants

There were 24 young participants who ranged from 17-24 years of age. There were also 24 elderly participants who ranged in age from 60-78 years.

Materials

21 familiar words and five unfamiliar words were selected from word norms (Toglia & Battig, 1978).

42 word sets were constructed for each participant; each word set consisted of randomly selected familiar and unfamiliar words (5-7 items for young participants and 4-6 items for elderly participants). There were no constraints on word selection for each set beyond random selection without replacement.

A prospective memory target was given to each participant depending on condition: participants in the familiar condition were given a familiar target word (i.e., "rake") while participants in the unfamiliar condition were given an unfamiliar target word (i.e. "sone"). None of the targets were in the 26 familiar/unfamiliar word sets.

Setup and Procedure

Participants were presented with 42 trials, each containing a word set as described above. All words for each trial were presented simultaneously for .75s per word. After the presentation of the words completed, the screen cleared and participants were instructed to recall and speak as many of the words in order as they could. The recall period was 1.5s per word

If their prospective memory target word showed up in their list of words, they were instructed to perform a keyboard action (hitting the "r" key). There were only 3 trials (one in each third) where the prospective memory target word appeared.

Measures

Einstein and McDaniel (1990) included several measures, but for our purposes we will focus on their retrospective memory for the items in each trial (labeled "Short term memory items" in the source article) and their prospective memory performance (labeled "Prospective memory" in the source article). The retrospective memory was the mean proportion of words recalled per trial and the prospective memory was the mean proportion of times they responded to the target word within that trial.

Results and Discussion

Empirical results are shown in Table 1. Unfortunately, no measures of variability were reported so they are unavailable.

As Table 1 suggests, younger adults recalled more items during the retrospective task than elderly adults, though there was no impact of the familiarity of the prospective memory target. Note that this difference emerged even though the elderly group had shorter average lists than the younger group.

In contrast, there was no difference between young and elderly on the overall prospective memory task. There was, however, an effect of the familiarity of the target: when the target was unfamiliar, performance on the prospective memory task was much better than when the target was familiar. Interestingly, there was no impact of age nor any interaction.

There are three main conclusions to draw from these results. First and unsurprisingly, the younger participants were better on the retrospective memory task than the elderly even though the elderly group had fewer items per trial. This overall finding is consistent with other research on short term memory in older populations (Luo & Craik, 2008; Rhodes et al., 2019).

Second, there was no difference in prospective memory for younger vs. elderly. While this finding is not uncontroversial (discussed in the general discussion), it has been replicated a number of times (Einstein & McDaniel, 1990; Cherry & LeCompte, 1999).

Finally, unfamiliar target words were easier to remember for a prospective memory task than familiar target words.

These findings suggest that different mechanisms account for the stark differences in performance between retrospective and prospective memory tasks. Specifically, we believe that a priming mechanism is a key component for successful prospective memory.

We next turn to the model architecture we use (ACT-R) and a description of the model to account for this data.

Architecture

A theoretical model was implemented to model the task performed by participants in the Einstein and McDaniel (1990) experiment. The model was developed using the ACT-R cognitive architecture (Anderson et al., 2004), a general theory of cognition that enables a simulation of an entire task (in this case, the experiment) while maintaining a high degree of psychological plausibility.

ACT-R consists of a number of modules, buffers, and a central pattern matcher. Modules in ACT-R contain a relatively specific cognitive faculty usually associated with a specific region of the brain. For each module, there are one or more buffers that communicate directly with that module as an interface to the rest of ACT-R. At any point in time, there may be at most one item in any individual buffer; thus, the module's job is to decide what and when to put a symbolic

Table 1: Proportion correct for both retrospective and prospective memory tasks (Einstein & McDaniel, 1990), exp. 2.

	Young		Elderly	
	Familiar	Unfamiliar	Familiar	Unfamiliar
Retrospective	.86	.82	.76	.80
Prospective	.28	.83	.36	.94

object into a buffer. The pattern matcher uses the contents of the buffer to match specific productions.

ACT-R uses if-then rules (productions) that will fire when their conditions are met by matching the contents of the buffers. If there is more than one production that can fire, the one with the highest utility (production strength) will fire. Each production can change either internal state (e.g., buffer contents) or perform an action (e.g., click a key). ACT-R interfaces with the outside world through the visual module, the aural module, the motor module, and the vocal module.

We describe the declarative memory system and the way that episodic goals are created here.

Declarative

The Declarative Module is the core method for memory retrievals within ACT-R including how a memory element (chunk) is encoded, remembered, and the rate at which where forgetting occurs.

Memory in ACT-R is described by a chunk's activation. Activation is the log odds that a particular chunk will be useful in the future; high activation chunks are expected to be very useful while low activation chunks are expected to be less useful.

Activation depends both on how much and how frequently a memory has been used in the past, as well as how related the item is to other memories that are currently the focus of attention. Activation consists of three primary components: activation strengthening, spreading activation, and noise. Activation strengthening is learned over time and is a function of how frequently and recently the memory has been thought about in the past, and represents the model's familiarity with a concept. Spreading activation is context dependent, allowing memories that are currently the focus of attention to activate, or prime, other related items. Noise is a random component added in to model the noise of the human brain. They are combined according to the following equation (Anderson, 2007):

$$A_i = B_i + \sum_j W_j S_{ji} + \varepsilon \tag{1}$$

where A_i is the total activation of chunk i, B_i is the total activation of chunk i, $W_j S_{ji}$ is activation spread from item j to item i, and ε is noise.

Spreading activation is spread along associations between memories. In addition to considering what items are being referenced at any given time, it also considers what items are in the current context. The current context consists of both those items being referenced, as well as the set of items in slot values of the items being referenced that are under consideration. Association strengths, intuitively, reflect how strongly item j, when currently being referenced, predicts that item i will be referenced next. The equations for the associative strength from an item j to an item i in memory are

$$S_{ji} = S - ln(fan_i) \tag{2}$$

where S_{ji} is the strength of association between chunks j and i, S is the maximum associative strength (mas parameter), and fan_j is the fan of chunk j (the number of other memory elements that memory j is associated with).

Episodic

The Episodic Module facilities the encoding and retrieval of episodic memories and is described more fully in Trafton, Altmann, and Ratwani (2011); it is not part of standard ACT-R but is an instantiation of the memory for goals theory (Altmann & Trafton, 2002). Not only can the episodic module be used for fine-grained cognitive control (Trafton et al., 2011; Altmann & Trafton, 2002; Altmann, Trafton, & Hambrick, 2017), but it can also be used to facilitate resumption after a task has been interrupted (Trafton et al., 2011; Trafton, Jacobs, & Harrison, 2012).

Episodic memories here are considered fine-grained snapshots of goals and events. In this model, episodic memories are created any time a new event occurs – in this case, every word on the screen is considered a new event, as is the start of each trial and the start of the experiment. When an episode is created, it receives a unique name to differentiate it from all other memories in declarative memory. It also contains a small bit of information about the episode itself (i.e., the experiment, the list, the room, etc.). These episodic markers can be used as keys to chain together other memories about events. The uniqueness of each episodic chunk also allows spreading activation to be especially potent because there will typically be far fewer links to a specific episode (Altmann & Trafton, 2002) than to a well-learned semantic memory. The episodic memory itself receives an increased starting level of activation (B_i in equation 1); this typically starts off as being rehearsed 3 times more than normal (Altmann & Trafton, 2002; Trafton et al., 2012). This increase in starting activation allows the episode to be relevant in the short term and allow more likely retrieval.

Prospective memory model

The prospective memory model emphasizes two processes: memory for goals to create episodic memories and resume the task after a prospective memory action has occurred; and elaboration to allow retrieval of prospective memories. As described below, both memory for goals and elaboration use spreading activation (described in the architecture section above) as a core process.

Memory for Goals

Memory for goals (Altmann & Trafton, 2002; Trafton et al., 2011) is a computational theory that describes how people construct, use, remember, and reconstruct goals. The theory suggests that episodic memories are created and strengthened when new goals and tasks are initiated. Priming allows the retrieval of these episodic memories when they are needed.

In the current model, a unique episodic chunk is created for each subtask of the experiment (i.e., each time a new list needs to be memorized) as well as for the prospective memory itself. These unique episodes have a functional component in this model: to serve a placekeeping function if the model gets interrupted and needs to resume.

Elaboration

We make the assumption that people elaborate information currently in working memory, remembering events, episodes, or information that may not be immediately relevant to their current goal (Kitajima & Polson, 1996). We also assume that elaborations are influenced by context through a spreading activation mechanism (Hiatt & Trafton, 2015; Taatgen et al., 2021).

Historical memories

The Einstein and McDaniel (1990) experiment relies on both familiar and unfamiliar words and the memory for goals and elaboration processes both use episodic memories. We modeled this by giving the model historical information for common and uncommon words and provided the model with a set of episodic memories.

The representation of both common and uncommon words was based on previous lexical words by (Reitter, Keller, & Moore, 2011) which included a lexical form and semantics. Concepts and meanings were separate chunks in memory that were linked to the lexical form via a slot in the chunk. Common words had very high initial activation (50 rehearsals) and many links to other concepts in memory (15 links). Uncommon words have very low initial activation and no links to other concepts in memory. Target prospective memory words had exactly the same representation, initial activation, and links to other chunks as their common or uncommon type (i.e., "rake" was a common word with a high activation and many links to other memories).

Episodic memories could contain task information (a task and a state marker) or an elaboration of a current memory linked to other memories through slot values. All episodic memories had a unique episodic cue so they would not become merged with other similar memories. Pre-existing episodic memories had a range of activation levels corresponding to 1-6 rehearsals.[1]

High level description of the model

In both the original experiment and the model, the prospective memory was told to the participants. The prospective memory was stored in declarative memory; an example from the familiar condition is shown below.

```
chunk429>
  isa elaboration
  memory-type episodic
  cue rake
  visual-cue "rake"
  action-type keyboard-action
  action "r"
  context psych-experiment
  episode episodic329
```

The model has two separate but related components: the retrospective list-memorization task and the prospective-memory task. Each of these components will be described separately, though they are interleaved in a single model.

Retrospective Memory Task. The retrospective memory component is similar to Anderson et al. (1998). When the model is presented with a list of words, it creates a unique identifier for this list (i.e., list42) that is used to link all the words it sees into a common list. The model also creates a unique episodic trace for the start of this goal (Altmann & Trafton, 2002; Trafton et al., 2011). The episodic trace will decay over time, but can be used to resume this task if it gets interrupted or the task gets suspended.

The model starts at the top of the screen and reads each word. After it reads a word, the model attempts to retrieve a meaning for that word. When a common word is displayed, the model remembers a meaning; when an uncommon word, the model cannot recall a meaning. Because the list of words was visible for a limited time, the model limited its time to remember a meaning. Regardless of whether the model was successful at remembering a word-meaning, the word was then put into declarative memory with a link for the unique list. The model then performed some elaboration on the meaning of the word itself, attempting to retrieve an episodic memory, similar to Hiatt and Trafton (2015) and Taatgen et al. (2021). Most of the time, an episodic memory was retrieved that may have been relevant to something else in declarative memory. This was also where a prospective memory could be retrieved; see below for further details.

After any elaboration occurred, the model progressed to the next word and the process restarted. When the trial time for that list was up, the screen cleared and the model attempted to

[1] All pre-existing historical knowledge is assuredly much larger in people, but higher values showed the same patterns described below.

remember the words on the list. The model used the unique list identifier as a source of spreading activation and spoke each word as it was remembered. When the recall stage of the experiment was completed and the time for the recall phase had completed, the next list was shown.

Prospective Memory Task. Recall that the model was told to strike a specific key when it saw a particular cue; this memory was stored in declarative memory as shown above.

When the model read a prospective memory target word and elaborated on it, the model had an opportunity to retrieve the prospective memory. If the prospective memory chunk was successfully retrieved, the model would verify that the target word was on the screen, notice it had a prospective memory action to execute, and perform the action. This prospective memory task changed the goal from the retrospective memory task, so after completion of the prospective memory action, the model needed to resume what it was doing (Monk, Boehm-Davis, Mason, & Trafton, 2004; Trafton & Monk, 2007; Trafton et al., 2011). One reason having a theoretical resumption process is so important is because the information for the unique list needed to be retrieved so it could be saved with any new words that appeared after the prospective memory cue.

In order to resume, the model attempted to recall the previous task-relevant episodic trace that it created at the beginning of the current list. If the model was successful, it was able to recall the specific goal, state, and list it was working on. If the model was unsuccessful, it needed to reconstruct the experimental context and try again. Under these conditions, the model was always able to recall the contextual information. The accuracy of each trial's retrospective and prospective tasks were recorded.

Success of the model during retrospective task. The model had quite high success on the retrospective task ($>$ 75%). This accuracy was achieved primarily through the spreading activation of the unique list. The model created a new experiment-chunk that contained the lexical name of the chunk, a link to its meaning, and a link to the unique list. During the recall phase, the unique list was a source of spreading activation and greatly facilitate retrieval. Note that the retrieved chunks were not the original historical chunks, but rather chunks created for this experiment and list. Base level activation played a modest but critical role in retrieval as well.

Failure of the model during retrospective task. Recall that every time a retrieval is made, transient noisewas added to the activation for every candidate chunk. Sometimes this noise causes the total activation to be below a retrieval threshold, causing a retrieval failure. Once a retrieval failure occurred, the model assumes there are no more words to remember, so completed the recall phase. Because transient noise is different for every chunk at every retrieval, the chance that activation for all candidate chunks is below the retrieval threshold increases as there are fewer words to remember.

Success of the model during prospective memory. When the model sees a prospective memory target, spreading activation flows to other related concepts. When the spreading activation is strong enough to boost the activation of the prospective memory chunk that was learned at the beginning of the experiment, the model successfully retrieves the prospective memory action rather than another episodic memory or elaboration. Note that for uncommon words, the fan is smaller so the impact of spreading activation is stronger.

Failure of the model during the prospective memory task. Sometimes the model sees the prospective memory target word but does not retrieve the prospective memory elaboration. There are several possible reasons. First, the amount of spreading activation to the target word target may not be very high; this is especially likely to occur for familiar words that have a high fan. Second, transient noise may cause another memory to have higher activation than the prospective memory elaboration, allowing another memory to be retrieved instead. Third, it is possible that transient noise has caused no elaborations to have sufficient activation to be retrieved; this is an extremely rare occurrence in this model since the historical memories and elaborations have a high activation.

Impact of age. As suggested earlier, the exact process of modeling age-related memory decline is still under debate. In our model, neither noise nor memory decay could account for the differences we see in the experimental data. Instead, we found, consistent with other theoretical work (Rhodes et al., 2019), that a deficit in memory retrieval accounted for the difference in age for retrospective memory. In ACT-R, the retrieval threshold is a parameter that controls what level of activation a chunk needs in order to be successfully retrieved and is a direct way to model deficits in memory retrieval.

Model Fit

Model fits were created by running the model 500 times for each condition, providing model stability. This model changed 3 variables from their defaults, retrieval threshold, activation source spread, and maximum associative strength. Parameter fits were obtained by searching through each parameter space. As discussed earlier, the retrieval threshold for the elderly model was changed to 1.5 from the default of 0 for the younger model. This was to model a deficit in memory retrieval that elderly seem to have (Rhodes et al., 2019). Model results showed the same qualitative results across most parameter values except in extreme circumstances (e.g., when retrieval threshold was extremely high and nothing could be recalled).

Source spreading activation for the imaginal buffer was set to 2.0, a common value. The maximum associative strength was set to a value of 3.1, a value that is a bit higher than normal. Base level activation was set at the default of .5 and activation noise was set at a traditional .1 value.

As Figure 1 suggests, the model is able to capture the three main effects of the empirical data. First, the elderly model has

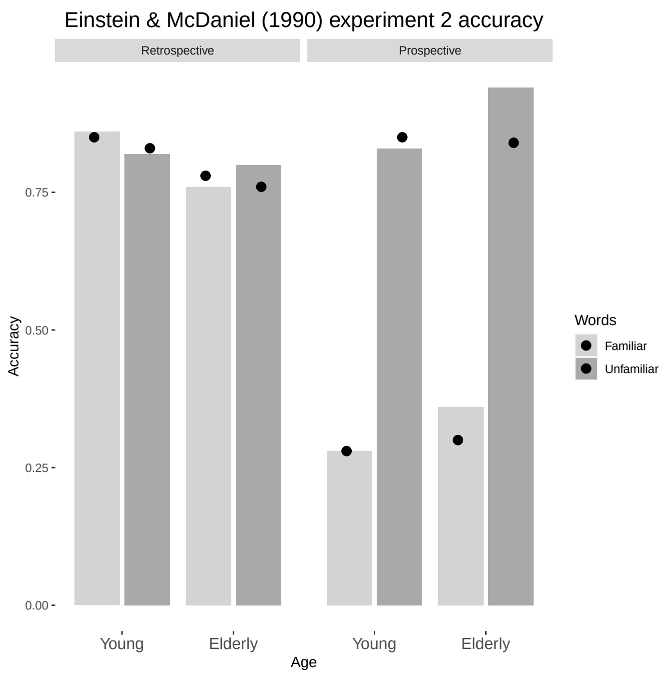

Figure 1: Accuracy for both retrospective and prospective memory tasks from Einstein and McDaniel (1990). Bars are empirical data and dots are model data.

worse retrospective memory than the younger model. Second, the model shows no difference in prospective memory for younger or older. Finally, when uncommon words were used as prospective memory cues, performance was far superior compared to common word cues.

The global shape of the model is quite similar to the empirical data, $R^2 = .97$. The model also fits very well in terms of its absolute deviation from the empirical data, RMSD = 0.045.

General Discussion

We described a high fidelity process model of prospective memory. Our model focuses on performance over reaction time. The model was created in the ACT-R cognitive architecture and uses two high-level mechanisms: memory for goals and elaboration. Memory for goals allowed the model to create episodic memories and then use those episodic memories to resume the primary task after a prospective memory had been completed. The elaboration mechanism allowed elaboration of words in working memory to occur which could then allow the recall of prospective memory goals based on environmental cues.

We solve the circular knowledge dependency problem (Li & Laird, 2013) by using a mix of procedural and declarative knowledge. We assume that there is general procedural knowledge that can execute when a prospective memory enters working memory (the retrieval buffer in our case). This prospective memory contains a cue that needs to be verified (a word) and an action to execute. Because the prospective memory uses declarative knowledge that can be created on the fly and it is general, the rule simply waits until the prospective memory gets retrieved. Our contextual priming provides an opportunity to retrieve the prospective elaboration when the relevant cue is perceived.

This model was fit to a classic (though controversial) empirical paper on prospective memory. The uncontroversial aspect of the model concerned the common and uncommon words; we used a standard spreading activation and fan effect mechanism to account for these findings.

The controversial aspect concerned the lack of difference between young and elderly participants. In fact, many researchers have attempted to conceptually replicate this finding with varying degrees of success and in fact meta-analysis suggests that prospective memory does decline with age, in contrast to the experiment modeled here (Uttl, 2008). While some researchers have highlighted experimental or analysis confounds (e.g., the small number of prospective memory items made it difficult to detect a significant difference), we believe that our model can provide some guidance as well.

One of the debates in the literature about prospective memory tasks focus on differences between laboratory and naturalistic tasks (Uttl, 2008). Our model suggests that regardless of the type of task, an important feature will be the size of the fan that exists for the prospective memory cue. If there is a large fan (which can occur in both experimental or naturalistic settings), prospective memory will be more difficult. Because naturalistic studies make it difficult to control the number of relationships, they are most likely to show a great deal of variability, which seems to be the empirical finding.

Another debate concerning prospective memory concerns the role of intelligence in prospective memory performance. The elderly participants in Einstein and McDaniel (1990) ended up having significantly higher WAIS-R scores than the younger participants (presumably because of the college environment the elderly were recruited from), which presents a possible confound in the empirical results. Other researchers have also found that IQ is related to prospective memory performance (Uttl, Graf, Miller, & Tuokko, 2001; Cherry & LeCompte, 1999; Maylor, 1996). Our model provides a straightforward explanation for this finding: people with a higher IQ should have a higher source activation (W in equation 1; Daily et al., 2001) or a higher pool for spreading activation (the S in equation 2).

Acknowledgments

This work was supported in part by ONR to GT. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the US Navy.

References

Altmann, E. M., & Trafton, J. G. (2002). Memory for goals: An activation-based model. *Cognitive science*, *26*(1), 39–83.

Altmann, E. M., Trafton, J. G., & Hambrick, D. Z. (2017). Effects of interruption length on procedural errors. *Journal of Experimental Psychology: Applied*, *23*(2), 216.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological review*, *111*(4), 1036.

Anderson, J. R., Bothell, D., Lebiere, C., & Matessa, M. (1998). An integrated theory of list memory. *Journal of Memory and Language*, *38*, 341-380.

Baddeley, A. (2012). Working memory: Theories, models, and controversies. *Annual review of psychology*, *63*, 1–29.

Cherry, K. E., & LeCompte, D. C. (1999). Age and individual differences influence prospective memory. *Psychology and Aging*, *14*(1), 60.

Daily, L. Z., Lovett, M. C., & Reder, L. M. (2001). Modeling individual differences in working memory performance: A source activation account. *Cognitive Science*, *25*(3), 315–353.

Einstein, G. O., & McDaniel, M. A. (1990). Normal aging and prospective memory. *Journal of Experimental Psychology: Learning, memory, and cognition*, *16*(4), 717-726.

Elio, R. (2006). On modeling intentions for prospective memory performance. In *Proceedings of the annual meeting of the cognitive science society.*

Ellis, J. (1996). Prospective memory or the realization of delayed intentions: A conceptual framework for research. *Prospective memory: Theory and applications*, 1–22.

Henson, R. N. (1998). Short-term memory for serial order: The start-end model. *Cognitive psychology*, *36*(2), 73–137.

Hiatt, L. M., & Trafton, J. G. (2015). A computational model of mind wandering. In *Proceedings of the annual meeting of the cognitive science society.*

Kitajima, M., & Polson, P. G. (1996). A comprehension-based model of exploration. In *Proceedings of the sigchi conference on human factors in computing systems* (pp. 324–331).

Laird, J. E. (2019). *The soar cognitive architecture*. MIT press.

Laird, J. E., Lebiere, C., & Rosenbloom, P. S. (2017). A standard model of the mind: Toward a common computational framework across artificial intelligence, cognitive science, neuroscience, and robotics. *Ai Magazine*, *38*(4), 13–26.

Lebiere, C., & Lee, F. J. (2002). Intention superiority effect: A context-switching account. *Cognitive Systems Research*, *3*(1), 57–65.

Li, J., & Laird, J. (2014). Spontaneous retrieval for prospective memory: Effects of encoding specificity and retention interval. In *Proceedings of the 29th aaai conference on artificial intelligence. aaai.*

Li, J., & Laird, J. E. (2013). The computational problem of prospective memory retrieval. In *Proceedings of the 12th international conference on cognitive modeling* (pp. 155–160).

Luo, L., & Craik, F. I. (2008). Aging and memory: A cognitive approach. *The Canadian Journal of Psychiatry*, *53*(6), 346–353.

Maylor, E. A. (1996). Age-related impairment in an event-based prospective-memory task. *Psychology and aging*, *11*(1), 74.

Monk, C. A., Boehm-Davis, D. A., Mason, G., & Trafton, J. G. (2004). Recovering from interruptions: Implications for driver distraction research. *Human factors*, *46*(4), 650–663.

Reitter, D., Keller, F., & Moore, J. D. (2011). A computational cognitive model of syntactic priming. *Cognitive science*, *35*(4), 587–637.

Rhodes, S., Greene, N. R., & Naveh-Benjamin, M. (2019). Age-related differences in recall and recognition: A meta-analysis. *Psychonomic Bulletin & Review*, *26*(5), 1529–1547.

Taatgen, N. A., van Vugt, M. K., Daamen, J., Katidioti, I., Huijser, S., & Borst, J. P. (2021). The resource-availability model of distraction and mind-wandering. *Cognitive Systems Research*, *68*, 84–104.

Toglia, M. P., & Battig, W. F. (1978). *Handbook of semantic word norms.* Lawrence Erlbaum.

Trafton, J. G., Altmann, E. M., & Ratwani, R. M. (2011). A memory for goals model of sequence errors. *Cognitive Systems Research*, *12*(2), 134–143.

Trafton, J. G., Hiatt, L. M., Harrison, A. M., Tamborello, F. P., Khemlani, S. S., & Schultz, A. C. (2013). Act-r/e: An embodied cognitive architecture for human-robot interaction. *Journal of Human-Robot Interaction*, *2*(1), 30–55.

Trafton, J. G., Jacobs, A., & Harrison, A. M. (2012). Building and verifying a predictive model of interruption resumption. *Proceedings of the IEEE*, *100*(3), 648–659.

Trafton, J. G., & Monk, C. A. (2007). Task interruptions. *Reviews of human factors and ergonomics*, *3*(1), 111–126.

Uttl, B. (2008). Transparent meta-analysis of prospective memory and aging. *PloS one*, *3*(2), e1568.

Uttl, B., Graf, P., Miller, J., & Tuokko, H. (2001). Pro- and retrospective memory in late adulthood. *Consciousness and Cognition*, *10*(4), 451–472.