

Towards a Computational Model of a Dynamic Feeling of Knowing

Justin Li (justinnhli@oxy.edu)
Bryce Boyle (bboyle@oxy.edu)
Occidental College, 1600 Campus Road
Los Angeles, CA 90041 USA

Abstract

Feelings of knowing (FOKs) are metamemory judgments that suggest an answer could be retrieved from memory with more effort. This paper reviews the psychological accounts of FOKs and maps them onto sources from the memory mechanisms of the common model of cognition. Two widely accepted accounts of FOK, that of cue familiarity and accessibility, map directly onto properties of the retrieval cue and result respectively. In considering these models of FOK, we identify an omission from the literature: the possibility that FOK changes over time while answering a question. We discuss the implications of this dynamic account and conclude with the difficulties of evaluating computational models of FOK.

Keywords: Feeling of Knowing; Metamemory; Knowledge Search; Common Model; Cognitive Architecture

Introduction

Feelings of knowing (FOKs) are a memory phenomenon where, despite not retrieving the answer to a question in the moment, people *feel* that they will be able to do so with more effort. FOKs have been studied as a topic of its own and as a way to gain insight into how memory is used in decision making (Nelson & Narens, 1994), with research focusing on the sources that underlie FOK and how it is influenced by the retrieval context. More recently, researchers have proposed the *cognitive-heuristic account* of metamemory: that FOKs serve the function of guiding memory retrieval, allowing for early failure without expending resources if the probability of finding an answer is low (Schwartz & Metcalfe, 2011). At the same time, although cognitive architectures such as ACT-R (Anderson, 2007) grew out of early models of declarative memory, metamemory phenomena has not received much attention and has not been the topic of cognitive modeling.

This paper complements the existing literature by exploring how feelings of knowing might be instantiated in a common model of cognitive architecture. We begin by placing FOKs in the context of knowledge search, and in doing so identify an omission from our current understanding of the phenomenon, that of how FOKs might change during question answering. We then briefly summarize a computational representation of memory, before committing the bulk of this paper to considering potential sources of FOK in the common model, as described by existing psychological theories. Returning to a dynamic FOK, we explore the consequences of such a theory, and conclude by discussing the obstacles that exist for evaluating a computational model of FOK.

FOK and Strategic Knowledge Search

In order to understand FOKs, its function within the memory systems of an agent must be understood: that of helping an

agent retrieve knowledge and engage in *knowledge search*. First proposed by Newell and Simon (1972), knowledge search is the process of finding knowledge that is relevant and could be applied to the current problem solving context. Newell and Simon do not elaborate on the processes of knowledge search, and knowledge search has received scant attention as compared to problem space search. Instead, the main advances come from psychological research on how people use memory in naturalistic settings, especially on how problem solving and decision making are intertwined with memory. The results show a rich landscape of memory processes: beyond basic recognition and recall, participants described determining recall specifications, gauging their own familiarity with the topic, relating multiple relevant memories, and verifying that a potential answer is in fact correct (Burgess & Shallice, 1996). All of these processes, together with actual memory retrievals, are necessary to answer one question. This account validates the idea that metamemory judgments such as FOK are used for the control of retrieval processes, such as selecting a search strategy and deciding whether to terminate search (Nelson & Narens, 1990). More recently, FOK has been shown experimentally to influence search termination and decision of what to rehearse (Singer & Tiede, 2008; Hanczakowski et al., 2014), further corroborating the cognitive-heuristic account of metamemory.

For example, consider a question such as *What film was nominated for seven Academy Awards in 1999?*¹ (Norman et al., 2016). To answer the question, an agent might retrieve cultural events in 1999, famous directors and actors/actresses, generally acclaimed films, and so on. Some of these results will be useful for answering the question; others may turn out to be irrelevant or lead to dead ends. In between these retrievals, FOKs play the role of determining the search strategy or whether to terminate search. Here, however, there is a mismatch between this hypothesized role of FOKs and how we conceptualize it for experiments. Thus far, experimental procedures for FOKs tend to only solicit a single judgment, either before or after the participants attempt to answer the question. If we accept that FOK is used to guide the multiple retrievals needed to find an answer, it raises the question: at which retrieval was the FOK solicited, and to which retrieval was the FOK indicating that an answer exists? When a participant reports their feeling of knowing, is it to the original question, or to any of the sub-questions that they ask themselves as they engage in the strategic search for the answer?

¹Answer: *Life is Beautiful*

Here we propose that the reported FOK is to the original question, and not to any of the other retrievals during the search process. This interpretation is more obvious for a question such as *What is the capital of Australia?*². Most people will suggest answers such as Sydney, Melbourne, and Brisbane before giving up. Despite these successful retrievals for Australian cities, participants will report that their FOK goes down over time before they terminate their search. That is, it seems clear to us that FOK is a *dynamic* signal that changes throughout the strategic memory search process: as additional retrievals are used for problem solving, the FOK fluctuates for the overarching goal of answering the original question. This is also consistent with the cognitive-heuristic account of FOK: in order for FOK to be a reliable signal for search termination, it *must* change over the course of the process to reflect whether an answer is still likely to be found. Again, this stands in contrast to how FOK is usually studied: all psychological experiments we have found only solicit participants’ FOKs once, either before or after they are given the chance to attempt to answer the question. While we have no doubt that such reports of FOK will still be correlated with the state of memory, ignoring the time course of FOK will likely omit crucial aspects of how the signal is determined. For the remainder of this paper, we will therefore assume this dynamic view of FOK as we consider how it might be modeled computationally.

Memory in the Common Model of Cognition

We now describe the agent framework in which we would like to model FOK, namely, that of the common model of cognition (Laird et al., 2017). The common model defines a set of representations and processes for modeling cognition, as implemented in cognitive architectures such as ACT-R and Soar (Anderson, 2007; Laird, 2012). Of particular interest to this paper are the declarative long-term memory (LTM) processes, specifically that of semantic memory, which we describe below.

Formally, the contents of LTM is an edge-labeled directed graph, defined by the tuple $\langle S, P, L, E \rangle$: S the set of entities or concepts (we use these terms interchangeably), which corresponds to the internal nodes of the graph; P the set of predicates, which corresponds to the edge labels of the graph; L the set of literals, such as numbers and strings, which corresponds to the leaf nodes of the graph; and E the set of direct edges from one entity to another entity or to a literal, $\langle s, p, o \rangle \in S \times P \times O$, with $O = S \cup L$. Borrowing from the knowledge representation literature, we will also refer to edges as *triples*, and refer to the elements of a triple $\langle s, p, o \rangle$ as the *subject*, the *predicate*, and the *object* respectively.

An agent has two ways of getting knowledge from LTM. First, for any entity s , the agent can *retrieve* all outgoing edges $\{\langle s, p, o \rangle \in E\}$ for which that entity is the subject. This mechanism is for accessing related information of a known concept, but to find an unknown concept that has certain

properties, the agent must *query* LTM instead. To do so, the agent creates a query *cue* $Q = \{q \in P \times O\}$, which describes the predicates and corresponding objects of the desired entity s such that $\forall \langle p, o \rangle \in Q, \langle s, p, o \rangle \in E$. We designate all matching entities of a query Q as S_Q , the set of retrieval *candidates*. If more than one such retrieval candidate exists, the entities with higher *base-level activation* are preferentially returned. Base-level activation is determined by $A(s) = \ln(\sum t_i^{-d})$, where t_i is the time since the entity s was last retrieved, and d is a decay rate parameter. Activation thus captures the recency and frequency of use of a concept, and is often used as a proxy of the importance of the concept to the agent at a particular time.

Within this framework, we can define the general process through which an FOK might be generated. When the agent is presented with a question, the agent would execute a sequence of queries and retrievals to LTM to attempt to answer the question. For clarity, we call the answer to the overarching question the *answer*, while an individual query will have a *result* (the entity that is returned) out of a set of *candidates* (other entities that match the cue). We assume that the FOK for the original question will change with each query and retrieval, and we are therefore interested in the computational processes that occur at those times and how they might affect the overall FOK.

A quick note on terminology: the term *retrieval* is overloaded in both psychology and cognitive architecture literature to sometimes mean both queries (with a cue) and retrievals (of a known concept in LTM). Retrieval will be used in the psychological sense in this paper; we will disambiguate the term as needed when talking about the specific computational mechanism.

Psychological Accounts of FOK

This section explores how psychological accounts of FOK might be realized within the common model. Within the psychology literature, there are three main accounts of FOK: cue familiarity, accessibility, and competition. For each, we first discuss the relevant psychological literature, before exploring how it may be translated computationally into long-term memory mechanisms. Since the literature primarily assumes a static FOK for a question, instead of one that changes over time, these computational models are all calculated from a single retrieval. A summary of these sources of information for FOK can be found in Table 1.

Two mathematical caveats must be considered. First, FOK may be a function of multiple parameters. Since we are primarily interested in what those parameters might be, and less interested in how they might be combined into a single FOK, we will assume that the function is a monotonic summary statistic denoted as $f(\cdot)$. Although the choice of summary statistic may affect the FOK calculation — the mean will be more sensitive to outliers than the median, for example — we consider this detail too low level for this paper. We do note that the competition account seems to be better modeled as the variance of a distribution than the mean or median, and

²Answer: Canberra

	Cue	Candidates	Result
Count		Accessibility	
1/Count		Competition	
Activation	Cue Familiarity	Competition	Accessibility
Connectivity	Cue Familiarity	Competition	Accessibility

Table 1: Difference sources of information for calculating FOK, and which corresponding psychology FOK theory they belong in. Gray cells represent sources that exist but not meaningful for FOK (e.g., the number of retrieval results, which is constant).

it is an open question whether or how other properties of the distribution might contribute to FOKs. Second, the domain of the output of the FOK function is unclear. The main constraint is that the agent should be able to determine whether an FOK judgment is high or low and thereby make retrieval decisions. The output could theoretically range over the real numbers — such as if FOK was the activation of a concept — with the agent learning decision thresholds over time. As we consider a dynamic FOK that may shift between difference sources of information, however, normalizing the FOK may be necessary, as the domain of the sources main differ wildly. For each account of FOK below, we will therefore also consider the population against which an FOK might be normalized.

The Cue Familiarity Account of FOK

As the name implies, the *cue familiarity* account of FOK focuses on the contribution of the retrieval cue to the feeling of knowing (Reder & Ritter, 1992; Metcalfe et al., 1993; Koriat & Levy-Sadot, 2001). For the purpose of this paper, we include all FOK sources that are based on the cue, including familiarity and domain knowledge (Schwartz, 1994). The intuition is that FOK is a summary of the amount of knowledge the agent might have about a topic, as estimated from the terms of the question. The more the agent is familiar or knowledgeable about the topic, the more likely that they will know the answer, leading to a higher FOK. Computationally, an FOK based on cue familiarity must be a function of the cue $Q = \{ \langle p, o \rangle \in P \times O \}$. In general, FOKs based on the cue familiarity account may be normalized against all concepts in LTM, as it would indicate the agent’s familiarity with these cues in particular, although care must be taken to account for cues that do not exist in LTM. We consider two metrics that might signify that the agent is “familiar” with the entities in the cue: their activations and their connectivities.

Activation One possible metric for the familiarity of the cue is the activation of each individual concept in the cue. Since activation reflects how recently and frequently a concept has been encountered, concepts with a high activation would be ones that are presented often, which in turn suggests that the agent would be familiar with them. Formally, this metric for FOK could be defined as:

$$\text{FOK} = \text{FOK}(Q) = f(A(o_1), \dots, A(o_n)) \forall \langle p, o \rangle \in Q$$

Connectivity In contrast to activation, connectivity captures how much knowledge the agent has of each concept in the

cue. A concept in which an agent is knowledgeable would be connected to many other concepts, while a concept of which the agent is ignorant would only be sparsely connected. In the extreme, the simple presence or absence of the concept (i.e., whether the agent recognizes the concept) may be a sufficient signal to terminate search, and it has been shown that recognition is can be a useful heuristic for knowledge search (Li et al., 2012).

The connectivity of a concept is measured by its *fan*, the number of incoming (*fan-in*) and outgoing (*fan-out*) edges. Arguments could be made for only considering fan-in or fan-out. The fan-in would represent the prevalence of the concept in different contexts, while the fan-out might represent its generality. It is also possible to consider the overall fan of a concept, regardless of the direction of the edges. More generally, connectivity may not just be the immediate neighbors of the cue, but the number of concepts within some neighborhood. We leave these implementation details as future work, and leave the precise meaning of the $\text{fan}(s \in S)$ function undefined. Formally, this metric for FOK could be defined as:

$$\text{FOK} = \text{FOK}(Q) = f(\text{fan}(o_1), \dots, \text{fan}(o_n)) \forall \langle p, o \rangle \in Q$$

The Accessibility Account of FOK

Unlike the cue familiarity account that depends on the cue, the accessibility account of FOK considers information that is only available during or after a retrieval, using the “byproducts” of the retrieval process (Koriat, 1993). The intuition behind the accessibility account is that the retrieval process may provide hints as to whether the agent could answer the question; if the first retrieval leads to a result with high confidence and certainty, this may lead to a high FOK even if additional retrievals are still necessary. Although the accessibility account includes uses of both properties of the result and metadata from the retrieval process, in practice the common model does not define universal a set of retrieval metadata that could be accessed. As a result, the models of FOK presented below are all functions of the retrieval result or the candidates.

While accessibility FOKs could also be normalized against other entities in LTM, a different reference group is also available: the set of candidates that matches the retrieval cue. This may reveal the relative importance of this result against other possible results. Such a comparison group would blur the difference between the accessibility account with the competition account, which we discuss in the next section.

Activation As with the activation of the cue, the activation of the result of a retrieval may be a metric for an accessibility-based FOK. Beyond summarizing the recency and frequency of use and therefore whether a concept is familiar, activation in this context may also represent the speed of the retrieval: the higher the activation, the more quickly the retrieval occurs. There is a large literature on the correlation between fluency and various memory phenomena (Alter & Oppenheimer, 2009), but here we consider it as equivalent to the activation of the retrieved result under the common model.

Formally, this metric for FOK could be defined as:

$$\text{FOK} = \text{FOK}(s) = f(A(s))$$

where s is the result of a retrieval.

Connectivity A different metric from the result of a retrieval is its connectivity, or the number of graph neighbors it has. As before, the connectivity of an entity represents the amount of knowledge that the agent has about the result. We note that the activation and connectivity of a retrieval result may be related due to spreading activation, but could also be the inverse of each other. A concept could be well-understood (i.e., have high connectivity) but irrelevant to the recent/current context (i.e., have low activation), as is the case when false memories are induced (Li & Kohanyi, 2016). Conversely, a concept could have low connectivity but high activation, such as when an agent is learning about a new concept.

Formally, this metric for FOK could be defined as:

$$\text{FOK} = \text{FOK}(s) = f(\text{fan}(s))$$

where s is the result of a retrieval.

Retrieval Candidates The accessibility account has an additional possible metric compared to cue familiarity: the number of candidates in the retrieval. The intuition for this metric is that if a retrieval cue matches many concepts, the agent might conclude that it has a lot of information at hand about the question, thus increasing the likelihood that it will be able to find the answer. Mathematically, this metric for FOK could be defined as:

$$\text{FOK} = \text{FOK}(s) = f(|S_Q|)$$

Other Accessibility Sources Other metadata of the retrieval process and the results have been proposed as FOK sources, although they do not map as cleanly onto the existing memory mechanisms of the common model. One such possibility is for FOK to be based on a *partial retrieval*, where some but not all information is retrieved (Hanczakowski et al., 2017). The intuition is that a partial retrieval suggests to the agent that a complete retrieval is possible, thus leading to an FOK. While this theory is psychologically plausible, we do not know of any common model cognitive architectures that support partial retrievals, leaving a model of such an FOK for future work.

Similarly, incorrect retrievals about the answer may contribute to FOK (Koriat, 1993). This source, however, may

be difficult to model computationally, as the agent has no a priori knowledge of whether a result is correct or not. The idea of incorrect retrievals as a source of FOK is further complicated by the idea that multiple retrievals are necessary to answer a question, as the majority of these intermediate results will not be the answer to the original question. On the other hand, this more complex landscape of knowledge search also presents opportunities. If “incorrect retrieval” is interpreted as the agent encountering difficulties, the need to change search strategies may itself decrease FOK, as it may suggest that the question is more difficult than assumed. More generally, it is not impossible for an FOK judgment to take other metacognitive phenomena into account. A thorough exploration of how FOK might relate to other metamemory is beyond the scope of this paper.

Finally, the semantic content of the result itself may be a source of FOK, allowing the agent to infer additional knowledge that boosts FOK. The inference process is highly dependent on the question and the existing knowledge and capabilities of the agent, however, and given the large space of possibilities, we leave the relationship between semantic content and FOK for future work.

The Competition Account of FOK

Less commonly discussed than the cue familiarity and accessibility accounts of FOK is the competition account (Schreiber & Nelson, 1998). Unlike the accessibility account where a large number of candidates suggest robust knowledge, the competition account states that FOK is *inversely* proportional to the number of candidates. While the competition account of FOK may technically be a subcategory of the accessibility account, in that the number of candidates is a piece of metadata from the retrieval process, we consider the competition account sufficiently different to address it separately. In particular, unlike the accessibility account where the result of a retrieval plays a main role, the competition account (in the extreme) does not consider the result at all. Rather, only the set of candidates influence FOK; which specific concept is retrieved is irrelevant.

The intuition for the competition account is that more potential results to a query increases the uncertainty as to which result is correct, thus decreasing FOK. A direct translation of the competition account is to use the inverse of the number of candidates, which could be defined as:

$$\text{FOK} = \text{FOK}(Q) = f\left(\frac{1}{|S_Q|}\right)$$

However, other metrics for the competition account is possible. Extending the idea of uncertainty caused by having many candidates, we could model competition using the distribution of the activation or connectivity of the candidates. A uniform distribution would indicate that no candidate is more likely than the other, suggesting uncertainty; in contrast, a peaked distribution would mean that the candidate with more probability mass is likely to be the correct answer. An

activation-based competition metric for FOK could be defined as:

$$\text{FOK} = \text{FOK}(Q) = f(\text{Var}(\{A(s) \forall s \in S_Q\}))$$

such that the larger the variance in activation, the larger the difference between the most activated concept and other concept, and therefore the more certain that it is the answer. As with other accounts, variance could be replaced with other summary statistics such as the interquartile range, as long as it correlated variance and inversely correlated with the uniformity of the activation values.

Hybrid Accounts

Although we have considered activation and connectivity as separate sources of FOK, cue familiarity and accessibility accounts of FOK could incorporate both sources of information. For example, FOK could be calculated by averaging the activation of neighboring concepts, resulting in an FOK that takes both activation and connectivity into account, combining more information from the agent’s knowledge base. Such a calculation is reminiscent of spreading activation, which bolsters its psychological plausibility. A systematic exploration of FOK metrics that combine sources, and their psychological plausibility, is beyond the scope of this paper.

Mixing and matching FOK accounts may apply at the higher level as well. While the cue familiarity and accessibility accounts each only take one type of memory metadata as input, in practice FOK may be the result of more complex combinations of these sources that together lead to the FOK that people report. This idea is not new, as it has been noted that cue familiarity is available after the question is asked but before a retrieval, while accessibility is only available during or after a retrieval. It has therefore been suggested that these could be used sequentially: that FOKs solicited earlier are a result of cue familiarity, and FOKs solicited later are a result of accessibility (Florer & Allen, 2000; Koriat & Levy-Sadot, 2001). These multiprocess theories hint at how the two accounts are not as independent as previously suggested. It is a small step from there to our proposed dynamic account of FOK, which we now turn our attention to.

A Dynamic FOK Account

As mentioned earlier in this paper, the psychology literature has focused on FOK as a single measurement during the process of question answering. We now return to the idea that FOK may instead be dynamic, changing over time as different strategies and multiple retrievals are used to answer a single question. The hybrid account of FOK hints at this possibility, by suggesting that FOK uses different sources depending on when it is solicited. One issue with this account, however, is the assumption that FOK is constant throughout answering a question, which would require different sources to somehow lead to the same resulting FOK. Instead, we reframe FOK as a judgment that is always changing while answering a question: different sources of FOK are used but do not have to agree with each other, and these fluctuations may in fact be part of how

FOK guides knowledge search. This hypothesis makes the our understanding of FOK more parsimonious, as the previously retrieved results (used by the accessibility account) are then used as cues for the next retrieval (used by the cue familiarity account), thus unifying the different accounts. This section considers the ramifications of this hypothesis, and proposes additional possibilities for how FOK may be determined.

First, we note that while results from past experiments are likely not invalid, they may only provide a narrow view of FOK. These measurements may only be accurate to the state of knowledge search at the time of solicitation, and without a detailed understanding of the memory search state of the participant, it is difficult to infer how the FOK was generated. Even assuming the cue familiarity or accessibility accounts, it raises questions as to what cues were used for familiarity judgments, or what retrieval metadata were used when accessibility was measured. The possibility of multiple retrievals that occur in sequence also muddle the distinction between retrieval cues and retrieval results, since the result of one retrieval may become the cue for the next retrieval. New experimental paradigms will need to be created to determine how FOK changes over time, before existing empirical can be integrated.

A dynamic FOK has implications not just for which sources are used (if they are indeed different sources at all), but what information each source provides. During the course of problem solving, the activation of entities will change based on the results of previous retrievals. A cue that initially had low activation may be boosted if multiple retrieval results are connected to it; conversely, previously highly activated entities may become less so over time. While the connectivity of LTM is less affected by retrievals, it is also not impossible that new connections could be made during problem solving, for example if an agent realizes that blue whales are not fish in answering *What is the largest fish on earth?*³ In sum, the sources do not only provide a single value, but a *history* of values which could be combined into an FOK judgment.

Access to a history of memory metadata raises the possibility that FOK could be based on *previous* FOK values, or at least some summary thereof. Consider again the question of what is the capital of Australia, and where an agent guess with several large Australian cities before giving up. This could be explained by the accessibility account using activation: more prominent cities such as Sydney are guessed first, before less-well-known cities like Perth, until the activation drops below some threshold and the agent terminates the search. However, another model of FOK is possible: that the search termination is not just due to the activation of the last retrieved concept, but due to the overall downward trend of activation. In this case, the FOK judgments are not based purely on activation, but is additionally modulated by how the FOK itself has changed over time. Mathematically, we might define FOK to be a function of time, FOK_t , with t being the number of steps in the past. In this example, the fact that $\text{FOK}_{-3} > \text{FOK}_{-2} > \text{FOK}_{-1}$

³Answer: Whale sharks

would further decrease the FOK judgment. More generally, FOK could be defined as

$$\text{FOK} = f(\text{FOK}_{-1}, \dots, \text{FOK}_{-T})$$

up to some time T in the past, plus additional inputs corresponding to the cue familiarity and accessibility accounts. Given the importance of history in this account, modeling FOK may therefore require a deep understanding of the landscape of memory processes and how they behave over time.

General Discussion

This paper has explored the possibilities for modeling feelings of knowing within the common model of cognition. The three main accounts of FOK — cue familiarity, accessibility, and competition — map well onto the existing architectural memory mechanisms. At the same time, the assumption that FOK is constant breaks down when multiple retrievals from long-term memory are needed to find an answer. As a result, we proposed the possibility of a dynamic FOK that changes over time as retrievals are made, and also raise the possibility that FOKs could take history into consideration.

Defining the mathematical space of FOK is a step forward, but evaluating computational models will be difficult. Matching human data may be possible if we restrict the model to questions that can be answered by a single retrieval. Experiments such as those reported in Schwartz et al. (2014) and Florer and Allen (2000) manipulate FOK by varying the amount of artificial context, thus creating new connections in LTM and also inducing unequal activation among the new concepts. Matching data on more complex questions, however, will be complicated by the multi-step retrieval process and uncertainty around which retrieval the FOK should be computed. Alternately, we can also foresee evaluation FOK on artificial agents, by examining which accounts most accurately predicts whether the agent is able to eventually find an answer. The advantage of this approach is that it can be applied on top of existing models of memory, and it may provide insight into how FOK may change over time. In the long run, cognitive models of FOK will have to meet both of these evaluation criteria in order to accurately reflect its function as a heuristic for memory retrieval in people.

References

- Alter, A. L., & Oppenheimer, D. M. (2009). Uniting the tribes of fluency to form a metacognitive nation. *Personality and Social Psychology Review, 13*(3).
- Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* Oxford University Press.
- Burgess, P. W., & Shallice, T. (1996). Confabulation and the control of recollection. *Memory, 4*(4).
- Florer, F. L., & Allen, G. (2000). Feelings of knowing in the Ranschburg effect. *The American Journal of Psychology, 113*(2).
- Hanczakowski, M., Zawadzka, K., & Cockcroft-McKay, C. (2014). Feeling of knowing and restudy choices. *Psych. Bullet. & Rev., 21*(6).
- Hanczakowski, M., Zawadzka, K., Collie, H., & Macken, B. (2017, 01). Metamemory in a familiar place: The effects of environmental context on feeling of knowing. *J. of Exp. Psych.: Learning, Memory, and Cognition, 43*(1).
- Koriat, A. (1993). How do we know that we know? The accessibility model of the feeling of knowing. *Psych. Rev., 100*(4).
- Koriat, A., & Levy-Sadot, R. (2001). The combined contributions of the cue-familiarity and accessibility heuristics to feelings of knowing. *J. of Exp. Psych.: Learning, Memory, and Cognition, 27*(1).
- Laird, J. E. (2012). *The Soar cognitive architecture*. MIT.
- Laird, J. E., Lebiere, C., & Rosenbloom, P. S. (2017). A standard model of the mind. *AI Magazine, 38*(4).
- Li, J., Derbinsky, N., & Laird, J. E. (2012). Functional interactions between memory and recognition judgments. In *Proceedings of the 26th AAAI conference*.
- Li, J., & Kohanyi, E. (2016). Towards modeling false memory with computational knowledge bases. In *Proceedings of the 14th ICCM*.
- Metcalfe, J., Schwartz, B. L., & Joaquim, S. G. (1993). The cue-familiarity heuristic in metacognition. *J. of Exp. Psych.: Learning, Memory, and Cognition, 19*(4).
- Nelson, T. O., & Narens, L. (1990). Metamemory: A theoretical framework and new findings. *Psychology of Learning and Motivation, 26*.
- Nelson, T. O., & Narens, L. (1994). Why investigate metacognition? In *Metacognition* (pp. 1–25). MIT.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Prentice Hall.
- Norman, E., Blakstad, O., Johnsen, O., Martinsen, S. K., & Price, M. C. (2016). The relationship between feelings-of-knowing and partial knowledge for general knowledge questions. *Frontiers in Psychology, 7*.
- Reder, L. M., & Ritter, F. E. (1992). What determines initial feeling of knowing? Familiarity with question terms, not with the answer. *J. of Exp. Psych.: Learning, Memory, and Cognition, 18*(3).
- Schreiber, T. A., & Nelson, D. L. (1998). The relation between feelings of knowing and the number of neighboring concepts linked to the test cue. *Memory & Cognition, 26*(5).
- Schwartz, B. L. (1994). Sources of information in metamemory: Judgments of learning and feelings of knowing. *Psych. Bullet. & Rev., 1*(3).
- Schwartz, B. L., & Metcalfe, J. (2011). Tip-of-the-tongue states: retrieval, behavior, and experience. *Memory & Cognition, 39*(5).
- Schwartz, B. L., Pillot, M., & Bacon, E. (2014). Contextual information influences the feeling of knowing in episodic memory. *Consciousness & Cognition, 29*.
- Singer, M., & Tiede, H. L. (2008). Feeling of knowing and duration of unsuccessful memory search. *Memory & Cognition, 36*(3).