

An Ecological Model of Memory and Inferences

Daniela Link (daniela.link@unil.ch) and Julian N. Marewski (julian.marewski@unil.ch)

Department of Organizational Behavior, Université de Lausanne, 1015 Lausanne, Switzerland

Lael J. Schooler (lschoole@syr.edu)

Department of Psychology, Syracuse University, Syracuse, NY 13244, USA

Abstract

In this paper, we develop a memory model that predicts retrieval characteristics of real-world facts. First, we show how ACT-R's memory model can be used to predict people's knowledge about real-world objects. The model assumes the probability of retrieving a chunk of information about an object and the time to retrieve this information depend on the pattern of prior environmental exposure to the object. Second, we use frequencies of information appearing on the Internet as a proxy for what information people would encounter in their natural environment, outside the laboratory. In two Experiments, we use this model to account for subjects' associative knowledge about real-world objects as well as the associated retrieval latencies. Third, in a computer simulation, we explore how such model predictions can be used to understand the workings and performance of decision strategies that operate on the contents of declarative memory.

Keywords: ACT-R; declarative memory; decision making; fast-and-frugal heuristics; Internet; strategy selection

The Importance of Memory for Inferences

Many of our every-day decisions are based on declarative knowledge retrieved from long-term memory. For example, a consumer who decides between different car brands will retrieve knowledge, such as information about the price segment, brand image, or fuel efficiency, to decide which brands to consider more closely. In judgment and decision research, there is a rich literature on how people infer unknown criteria, such as the quality of a car, based on object attributes used as *cues* (e.g., Gigerenzer, Hertwig, & Pachur, 2011). The kind of *cue-knowledge* a person retrieves when making a decision will likely depend on the information related to the decision objects she has encountered before, say, on the Internet. The person will then use this cue-knowledge as input of decision strategies when making inferences about unknown *criteria*, such as the quality of a car. A detailed cognitive model of how environmental patterns are reflected in memory, tied to models of decision making, could hence help to understand how human decision making depends on and is adapted to the environment.

Modeling Declarative Knowledge in ACT-R

In the cognitive architecture *ACT-R*, knowledge about the world is represented in declarative memory (Anderson, Bothell, Byrne, Douglass, Lebiere, & Qin, 2004). The basic unit of knowledge in declarative memory is the *chunk*. A chunk combines a set of elements into a long-term memory

unit, where different concepts are configured together in the chunk's *slots*. New declarative knowledge is added to memory by encoding representations of objects that are attended in the environment. For example, the knowledge that the city of Berlin has an airport can be represented in a chunk with the following structure:

```

BERLIN-AIRPORT
ISA           CITY_FACT
CITY         BERLIN
ATTRIBUTE    AIRPORT
RELATION     HAS
    
```

The chunk is of type *CITY_FACT*. Its slots contain the city *BERLIN*, the attribute *AIRPORT*, and the relation *HAS*. If the same constellation of concepts is encountered and attended again, rather than creating a duplicate chunk, the memory entry of an existing chunk will be strengthened, and as a result, will become more readily accessible in memory.

In addition to symbolic information (the chunk-type and slot values), each chunk encodes subsymbolic information about the likelihood that the chunk will be needed to reach one of the system's processing goals – the chunk's *activation* (Anderson & Milson, 1989). A chunk's activation, in turn, is probabilistically related to its retrieval and the time required for retrieval. Table 1 summarizes the relevant equations for ACT-R's declarative memory system (see Anderson et al., 2004 for details on ACT-R's theory of declarative memory).

Table 1: Equations relevant for memory retrieval

Equation number	Equation
1) Activation	$A_i = B_i + \sum_{j=1}^m W_j S_{ji}$
2) Base-level Activation	$B_i = \ln n / (1-d) - d \ln L$
3) Associative Strength	$S_{ji} = \ln \frac{P(i j)}{P(i)}$
4) Attention Weighting	$W_j = \frac{W}{m}$
5) Retrieval Probability	$p_i = \frac{1}{1 + e^{-(A_i - \tau)/s}}$
6) Retrieval Time	$T_i = F e^{-A_i}$

Note. Equation 2 is an approximation of base-level activation.

In the following, we explore to what extent ACT-R's memory model can be used to predict people's knowledge about real-world objects when using the Internet as a mirror

of the environment. Implementing the above formulas in Matlab, we aspire to develop a convenient method for modeling memory contingent on frequencies of objects and their attributes in people’s natural environment, outside the laboratory. As we will illustrate, alike efforts may be helpful when using ACT-R’s memory system to understand, for instance, the workings and performance of decision strategies.

Predicting Retrieval from Internet Frequencies

We used frequencies of mentions of objects on the Internet as proxies to what information people would encounter in their natural environment, outside the laboratory (henceforth *web frequency*). Specifically, we searched for the names of cities (e.g., Berlin), the names of city attributes (e.g., airport) and the combination of cities and attributes (e.g., Berlin airport) on www.en.wikipedia.org in English, using the Wikipedia API tool to find the total number of hits for pages containing our search term. Search was performed on September 14th 2014.

To calibrate and test our model, in two Experiments, we collected behavioral data on people’s knowledge of European cities by asking them for pairs of cities and attributes whether or not they had heard of the city having the attribute before (e.g., “Does Berlin have an airport?”).

Memory Activation of Knowledge

We assume that when a person believes she has encountered a city together with an attribute before, this implies a successful retrieval of a chunk i encoding the relation between the city and the attribute. Given this assumption, each time a person tries to retrieve a combination of a city and an attribute, according to Equation 1, two sources of activation for chunk i are (a) the *base-level activation* of chunk i whose slots contain the city name and the attribute name and (b) *spreading activation* from the city and attribute names included in the retrieval request.

A chunk’s base-level activation is a function of the number of encounters, n , with the object or relation encoded in the chunk (Equation 2). We approximate the frequency with which a city and an attribute occur together in one context by the number of times both concepts co-occur on one page in the knowledge base Wikipedia ($N_{\text{city}\&\text{cue}}$). As a simplification, we assume the time of creation to be equal for all chunks and hence replace the lifetime of the chunk, L , by a constant. The decay parameter d is usually set to .5. The base-level activation of the chunk i can hence be written as:

$$B_{i,\text{web}} = c_0 + \ln N_{\text{city}\&\text{cue}}, \quad (7)$$

where the constant c_0 absorbs the value for the term $\ln 2 - .5 \ln L$.

For the spreading activation, we assume that the chunks encoding the city and the attribute forming part of the retrieval request spread activation to chunk i . Following Equation 3, the associative strength between a city and an attribute depends on the number of times the city and the

attribute co-occur together relative to each of the concepts’ base-rates of occurring (Schooler & Anderson, 1997). $P(\text{city}|\text{cue})$ is the probability that the city occurs, given the presence of the attribute. When dividing this conditional probability by $P(\text{city})$, we get a measure for how much more likely the city is to occur when the attribute is present than when it is not. The associative strength between the attribute and the city is the logarithm of this odds ratio:

$$S_{\text{cue},i} = \ln \frac{P(\text{city}|\text{cue})}{P(\text{city})} = \ln \frac{\frac{P(\text{city} \cap \text{cue})}{P(\text{cue})}}{P(\text{city})} = \ln \frac{P(\text{city} \cap \text{cue})}{P(\text{city}) * P(\text{cue})}. \quad (8)$$

We estimate the probability of a city or attribute occurring, ($P(\text{city})$, $P(\text{cue})$) by dividing the frequency of its occurrence in the knowledge base (N_{city} , N_{cue}) by the total size of the knowledge base N_{all} . We approximate the size of the knowledge base by the number of hits for pages in Wikipedia on which the word *and* appears (N_{and}), so we can write:

$$S_{\text{cue},i} = \ln \frac{\frac{N_{\text{city}\&\text{cue}}}{N_{\text{and}}}}{\frac{N_{\text{city}}}{N_{\text{and}}} * \frac{N_{\text{cue}}}{N_{\text{and}}}} = \ln \frac{N_{\text{city}\&\text{cue}} N_{\text{and}}}{N_{\text{city}} N_{\text{cue}}}. \quad (9)$$

It can be shown mathematically that $S_{\text{city},i} = S_{\text{cue},i}$. Assuming that the attention weights W_j from the m sources of activation sum to 1 (cf., Anderson, Reder, & Lebiere, 1996) and activation spreading from the city and the attribute with equal proportions, the total activation for chunk i , as estimated from the web, can be written as:

$$A_{i,\text{web}} = c_0 + \ln N_{\text{city}\&\text{cue}} + \ln \frac{N_{\text{city}\&\text{cue}} N_{\text{and}}}{N_{\text{city}} N_{\text{cue}}}. \quad (10)$$

We assume the memory activation A_i resulting from encounters with information in a person’s environment to be a linear function of the activation $A_{i,\text{web}}$ as estimated from web frequencies:

$$A_i = c + b A_{i,\text{web}}. \quad (11)$$

The parameters c and b serve as scaling parameters describing the unknown relation between how often a person encounters an object in her environment and the web frequency of the corresponding search term.

Given these assumptions, the formula approximating memory activation for chunk i by web frequencies of corresponding search terms N can be written as:

$$A_i = c + b \left(\ln N_{\text{city}\&\text{cue}} + \ln \frac{N_{\text{city}\&\text{cue}} N_{\text{and}}}{N_{\text{city}} N_{\text{cue}}} \right). \quad (12)$$

Retrieval Probability & Retrieval Latency of Knowledge

We use the chunk’s activation estimated from web frequencies to predict our participants’ *retrieval probabilities* of city-attribute associations according to Equation 5:

$$p_i = \frac{1}{1 + e^{\frac{-(c+b) \left(\ln N_{\text{city}\&\text{cue}} + \ln \frac{N_{\text{city}\&\text{cue}} N_{\text{and}}}{N_{\text{city}} N_{\text{cue}}} \right) - \tau}{s}}}, \quad (13)$$

as well as corresponding *retrieval times* according to Equation 6:

$$T_i = F e^{-\left(c+b \left(\ln N_{\text{city}\&\text{cue}} + \ln \frac{N_{\text{city}\&\text{cue}} N_{\text{and}}}{N_{\text{city}} N_{\text{cue}}}\right) + \varepsilon\right)} \quad (14)$$

where τ is the *retrieval threshold* and ε is the *retrieval noise* generated from a logistic distribution with a mean of zero and a variance of $\sigma^2 = \frac{\tau^2}{3} s^2$. In our model, we assume noise not only in the activation level but also in the retrieval threshold (cf., Marewski & Schooler, 2011), where the total retrieval noise parameter, s , is an aggregate of the *criterion noise parameter*, s_τ , and *activation noise parameter*, s_A , so that

$$s = \sqrt{s_\tau^2 + s_A^2}. \quad (15)$$

The response times are assumed to be the sum of *perceptual-motor times*, I , such as visual encoding and key pressing, and memory retrieval time:

$$RT_i = I + T_i. \quad (16)$$

Empirical Data

Participants

One hundred twenty-eight (Exp. 1) and 73 subjects (Exp. 2) participated in an experiment conducted at the University of Lausanne, Switzerland. Participants received a flat fee of CHF 5 (\$ 5.14) plus a bonus of up to CHF 33 (\$ 33.90) depending on the consistency of their responses in the main task and a short control task at the end of the experiment.

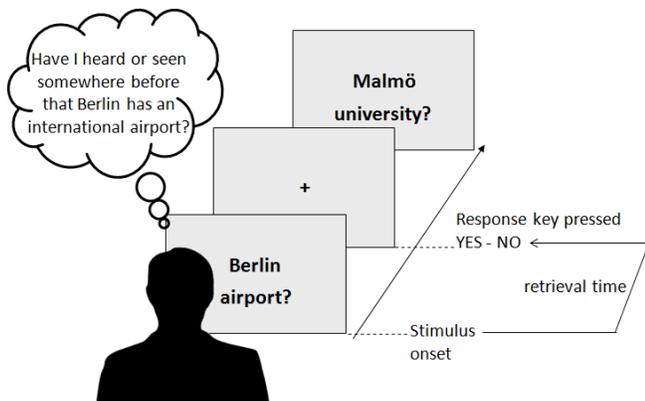


Figure 1: Illustration of the city-attribute task. Attributes were *international airport*, *university*, *high-speed train line*, *major league soccer team*, *company*, *underground*, *cathedral*, and *harbor* (Exp. 1).

Stimuli and Procedure

We assessed retrieval rates and response time distributions for people’s knowledge about 95 European cities on 8 attributes in Experiment 1 and 7 in Experiment 2. The difference between Experiment 1 and Experiment 2 was that we dropped the attribute *harbor* from the list of attributes for which knowledge was tested. Specifically, participants saw city-attribute pairs, one at a time (Figure 1). Participants were asked to respond by key press either with “Yes” (they could remember having heard or seen somewhere before

that the city possessed the attribute) or with “No” (they could not recall any such instances). For each trial, we recorded both (i) subjects’ responses and (ii) the time that elapsed between the presentation of a city-attribute pair and a keystroke. Additionally, for each city, we asked subjects whether they recognized the city name and whether they knew anything else about the city. In total, subjects responded to 950 (Exp. 1) or 855 (Exp. 2) trials.

Model Fits and Predictions

We fitted the parameters of the memory model for chunks encoding knowledge about cities to the data from Experiment 1. Leaving these parameters fixed, we used our model to predict (i.e., for a different set of participants) memory performance in Experiment 2.

Model Calibration on Experiment 1

Post-hoc, the cities “Nice” and “Derby” were excluded because web frequencies also included results for the adjective “nice” and the sport “derby”. Also all Swiss cities were excluded from the list because knowledge about these cities reflected personal experience rather than knowledge acquired through the media. 88 cities were included in the final sample. To calibrate the model, we first fit Equation 13 to the observed retrieval rates from Experiment 1. We set the total retrieval noise s to the value (.83) used by Marewski and Schooler (2011) and anchored the activation scale by setting the expected value of the retrieval criterion distribution, τ , to zero, so that an object with an activation of 0 would have a 50% chance of being retrieved (cf. Marewski & Schooler, 2011). With a simple regression conducted on the log-odds form of Equation 13, we estimated the constant c (-6.11) and the scaling parameter b (.69). The Pearson correlation between empirical retrieval rates and simulated retrieval probabilities is $r = .72$.

With these parameters fixed, in a second calibration step, we fit Equation 16 to the response time distributions for successful retrievals (“Yes” responses) in Experiment 1. Of course, response latency is not a perfect proxy for retrieval time. The total response time includes other components such as the time it takes to read a word and the time to press a key. To model these non-retrieval times, we assume response times are the sum of retrieval times plus perceptual-motor times (Equation 16). We model perceptual-motor times by drawing from a uniform distribution with boundaries of $t \pm t/2$, where t is set to the mean time as simulated by the ACT-R production rules necessary for performing the task excluding the memory retrieval (1.01 s).

Subsequently, we fit the latency factor F (.80) and the criterion noise parameter s_τ (.67) to the response time distributions of the 704 items of the city-attribute task in Experiment 1. This, as implied by Equation 15, fixes the activation noise parameter to $s_A = .49$. We did so by minimizing the sum of maximum vertical distances between the empirical and predicted cumulative response time distributions (cf., Voss, Rothermund, & Voss, 2004),

weighted by the empirical retrieval rates for each of the items. The simulation calculates, for each item, the expected value of retrieval time from the proportion of an item’s activation distribution that falls above a retrieval criterion sampled from the retrieval criterion distribution. To simulate response time distributions, we took a total of 10,000 samples from the retrieval criterion distribution per item. In sum, four parameters were estimated to predict retrieval probability and retrieval latency. In addition, four parameters were fixed (i.e., not fitted to the data). Table 2 gives an overview of the parameters and their values.

Table 2: Parameters of the memory model

Parameter	Value
Parameters estimated from retrieval probabilities	
Constant, c	-6.11
Scaling parameter, b	.69
Parameters estimated from response time distributions	
Latency factor, F	.80
Criterion noise parameter, s_τ	.67
Fixed parameters	
Total retrieval noise parameter, s	.83
Expected value of retrieval criterion distribution, τ	0
Activation noise parameter, s_A	.49
Expected value of the perceptual-motor time distribution	1.01 s

We calculated medians of the empirical and simulated response time distributions excluding city-attribute pairs for which less than three participants responded with “Yes”. We then smoothed the empirical and simulated medians with a running window of size five. The weighted (by the number of “Yes” responses) correlation between empirical and simulated smoothed median response times is $r = .62$.

Model Predictions for Experiment 2

Leaving these parameter values unchanged, we predict memory performance in Experiment 2. Figures 2 and 3 show the predicted and observed retrieval rates and response time distributions, respectively.

Figure 2 plots retrieval as a function of activation. The points represent the empirical retrieval rates (proportion of “Yes” responses), the S-shaped curve shows the predicted retrieval probabilities based on Equation 13. The Pearson correlation between empirical retrieval rates and predicted retrieval probabilities is $r = .72$.

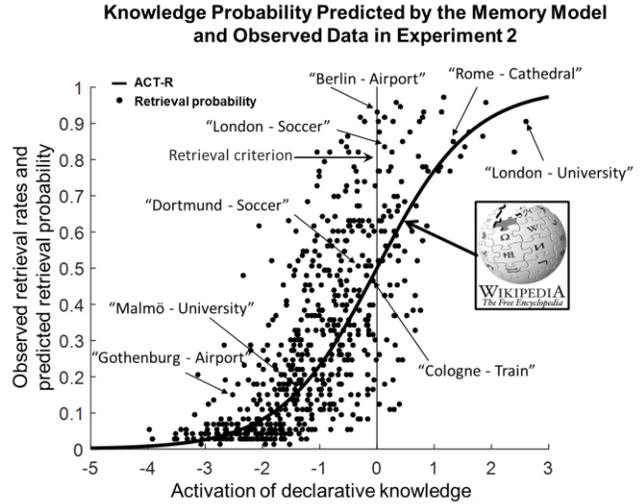


Figure 2: Observed retrieval rates and predicted retrieval probabilities for knowledge about 88 cities (Exp. 2) computed over 73 participants. Retrieval rates are plotted as a function of the expected value of the knowledge activations for 616 city-attribute pairs. The vertical line shows the expected value of the retrieval criterion.

Figure 3 plots response times for positive responses (“Yes”) given to the city-attribute task of Experiment 2 as a function of the corresponding chunk’s expected value of activation. The points represent the empirical quartiles of response time distributions, the solid lines show the quartiles of predicted response time distributions based on ACT-R’s retrieval mechanism (Equation 14). As can be seen, while generally increasing with decreasing activation, median response times are not a simple monotonic function of a chunk’s expected value of activation. Chunks will be retrieved when their momentary activation exceeds the retrieval threshold. As we assume noise in a chunk’s activation as well as in the retrieval criterion, chunks with a low expected value of activation sometimes exceed the retrieval criterion, at a momentary activation that is likely higher than the expected value of their activation. For that reason, predicted response times flatten out towards the lower end of the activation scale. As Figure 3 shows, our memory model is able to capture the increase in median and spread of response time distributions with decreasing activation of memory chunks. Response time distributions based on a low number of “Yes” responses are noisier and less well predicted by our memory model than those calculated from a high number of responses. Excluding city-attribute pairs for which less than three participants responded “Yes”, the weighted (by the number of “Yes” responses) correlation between empirical and predicted smoothed median response times is $r = .34$.

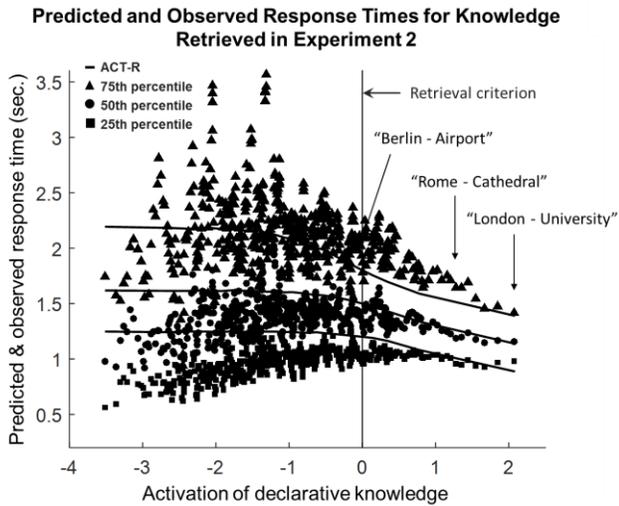


Figure 3: Dots show the 25th, 50th, and 75th percentiles of measured response times for the 572 city-attribute pairs of Experiment 2 where at least three participants responded with “Yes” in the city-attribute task. The solid lines show the quartiles of response times predicted by the model. The x-axis plots the expected value of the chunk’s activation, as derived from web frequencies. Response times and activations are smoothed with a moving average. The vertical line shows the expected value of the retrieval criterion.

Application to Decision Making

We extend previous efforts to populate the contents of ACT-R’s declarative memory with records that reflect the associated objects’ statistical patterns of occurrence in the real world (Marewski & Schooler, 2011; Salvucci, 2014). We believe modeling efforts of this kind lend themselves to many possible applications. In what follows, we illustrate just one.

Much research in the cognitive and decision sciences has explored how people infer objects’ (e.g., cities’) values on unknown criteria (e.g., population size, wealth) from the objects’ attributes, used as cues. Within the *fast-and-frugal heuristics* research program (e.g., Gigerenzer et al., 2011) several strategies describing how people make such inferential decisions have been suggested. The *take-the-best heuristic* (Gigerenzer & Goldstein, 1996), for example, retrieves knowledge about objects’ attributes in order of strength of relation to the criterion. This strength of relation is measured as *cue validity*, or as the probability that a city A has a higher value on the criterion (e.g., population) to be inferred than city B, given that A has a positive value on a cue (e.g., has a university) and B a negative or unknown value on that cue (e.g., has no university). Starting with the most valid cue, take-the-best prescribes comparing objects successively on cues in order of decreasing validity, until one cue is identified that allows for making a decision.

While there is evidence that people actually use strategies like take-the-best (e.g., Bröder & Gaissmaier, 2007; Walsh

& Gluck, 2016), and are able to adapt their strategy choice to the statistical structure of the environment (e.g., Rieskamp & Otto, 2006), what is known as the *strategy selection problem* remains a serious modeling challenge in the cognitive decision science and beyond (see Marewski & Link, 2014 for an overview).

In addressing that modeling challenge, one strand of research explores how environments are reflected in the memory system, that is, how statistical properties of the environment translate into retrieval probabilities and retrieval latencies of decision-relevant information. In interaction with the memory system, so the rationale goes, the environment carves out for each strategy a *cognitive niche* (Marewski & Schooler, 2011). In so doing, that interplay likely restricts the consideration set of strategies that can be applied to make a decision. Second, among the set of *applicable* strategies, currencies like the strategies’ speed of execution, required effort, and accuracy influence selection.

The memory model introduced in this paper simulates which knowledge a person will likely retrieve when confronted with a decision problem. In doing so, the model generates knowledge which can serve as input for different decision strategies. Given the rules prescribed by a particular strategy, one can make predictions on how a strategy will operate, based on the input provided by the memory model. In this way, the model aids exploring whether a strategy will be applicable, how much effort executing that strategy will require (e.g., the number of cues that must be retrieved before a decision can be made), and how accurate the resulting decisions might be.

To illustrate this, Figure 4 explores the niche of the take-the-best heuristic: Panel A depicts the probability of applicability of this heuristic, B the mean cue validity of the discriminating cue, and C the mean accuracy across paired comparisons of 88 cities included in Experiments 1 and 2. Cue validities were calculated from the actual attributes of the cities for a comparison of city size. The probability of attribute-knowledge retrieval was simulated based on the memory model calibrated to the retrieval rates observed in Experiment 1. The cities have been grouped into 22 equally sized bins according to their rank of environmental frequency (approximated by web frequencies).

As can be seen in Panel A, the probability that take-the-best can make a decision increases with the environmental frequencies of both cities. This relationship is paralleled by the effort required to execute take-the-best (Panel B): Fewer cues need to be checked (i.e., the discriminating cue is of high validity) as the environmental frequencies of the cities increase. In areas where both cities are of low environmental frequency, the applicability of take-the-best is at its lowest, and in the cases where that heuristic is applicable, it needs to examine several cues before a decision can be made. As one might expect, the heuristic’s accuracy (Panel C) generally rises with the validity of the discriminating cue. However, accuracy is low when both cities have about the same environmental frequency.

Outlook and Conclusion

We are working on implementing simulations of memory-based inferences to, eventually, predict when people will use which decision strategy in a given environment. We hope that such modeling efforts will, one day, invite insights into how the environment, in interaction with the memory system, aids adaptive strategy selection.

Acknowledgments

This work was supported by a grant from the Swiss National Science Foundation [100014_144413].

References

- Anderson, J. R., & Milson, R. (1989). Human memory: An adaptive perspective. *Psychological Review*, *96*, 703–719.
- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, *111*, 1036–1060.
- Anderson, J. R., Reder, L. M., & Lebiere, C. (1996). Working memory: Activation limitations on retrieval. *Cognitive Psychology*, *30*, 221–256.
- Bröder, A., & Gaissmaier, W. (2007). Sequential processing of cues in memory-based multi-attribute decisions. *Psychonomic Bulletin & Review*, *14*, 895–900.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *103*, 650–669.
- Gigerenzer, G., Hertwig, R., & Pachur, T. (Eds.). (2011). *Heuristics: The foundations of adaptive behavior*. New York: Oxford University Press.
- Marewski, J. N., & Link, D. (2014). Strategy selection: An introduction to the modeling challenge. *Wiley Interdisciplinary Reviews: Cognitive Science*, *5*, 39–59.
- Marewski, J. N., & Schooler, L. J. (2011). Cognitive niches: An ecological model of strategy selection. *Psychological Review*, *118*, 393–437.
- Rieskamp, J., & Otto, P. E. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, *135*, 207–236.
- Salvucci, D. D. (2014). Endowing a cognitive architecture with world knowledge. In P. Bello, M. Guarini, M. McShane, & B. Scassellati (Eds.), *Proceedings of the 36th Annual Conference of the Cognitive Science Society* (pp. 1353-1358). Austin, TX: Cognitive Science Society.
- Schooler, L. J., & Anderson, J. R. (1997). The role of process in the rational analysis of memory. *Cognitive Psychology*, *32*, 219–250.
- Voss, A., Rothermund, K., & Voss, J. (2004). Interpreting the parameters of the diffusion model: An empirical validation. *Memory & Cognition*, *32*, 1206-1220.
- Walsh, M. M., & Gluck, K. A. (2016). Verbalization of decision strategies in multiple-cue probabilistic inference. *Journal of Behavioral Decision Making*, *29*, 78-91.

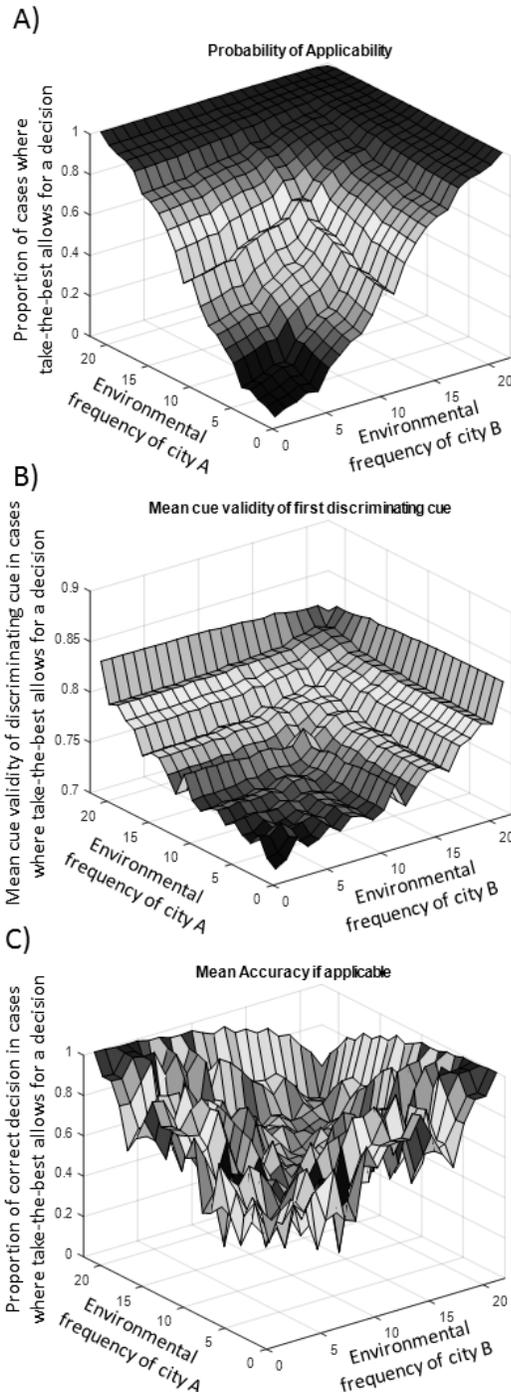


Figure 4: Simulation of inferences about city size made by the take-the-best heuristic. The 88 cities are grouped into 22 equally-sized bins according to their rank of environmental frequency. Bin numbers are shown on the horizontal axes. The vertical axis shows the mean applicability (A), cue validity of the first discriminating cue (B) and accuracy of inferences (C) across 10,000 simulated subjects for an exhaustive pairing of cities within each of the bins. Note that these simulations are exploratory.