

Recommender Systems for Literature Selection: A Competition between Decision Making and Memory Models

Leendert van Maanen (leendert@ai.rug.nl)

Department of Artificial Intelligence, University of Groningen
P.O.Box 407 9700 AK Groningen, The Netherlands

Julian N. Marewski (marewski@mpib-berlin.mpg.de)

Max Planck Institute for Human Development, Center for Adaptive Behavior and Cognition
Lentzeallee 94, 14195 Berlin, Germany

Abstract

We examine the ability of five cognitive models to predict what publications scientists decide to read. The cognitive models are (i) the Publication Assistant, a literature recommender system that is based on a rational analysis of memory and the ACT-R cognitive architecture; (ii-iv) three simple decision heuristics, including two lexicographic ones called *take-the-best* and *naïveLex*, as well as unit-weight linear model, and (v) a more complex weighted-additive decision strategy called Franklin's rule. In an experiment with scientists as participants, we pit these models against (vi) multiple regression. Among the cognitive models, *take-the-best* predicts most scientists' literature preferences best. Altogether, the study shows that individual differences in scientific literature selection may be accounted for by different decision-making strategies.

Keywords: Recommender system; ACT-R; rational analysis; simple heuristics; *take-the-best*; literature search

Literature Selection

In 2006, the number of scientific publications in the relatively small ISI subject category *Information Science & Library Science* was 2054. In other words, researchers working in this area had to scan through over 2,000 papers a year to keep up with the current developments. However, this number is, if anything, an underestimation of the total number of potentially relevant papers, as this number only holds if a researcher is interested in a single subject area. In practice, most researchers work on the intersection of multiple domains, increasing the number of potentially relevant papers enormously. Not only professionals in the scientific domain are confronted with masses of potentially relevant information. Also, government or business employees often need to decide which of numerous reports, leaflets, and bulletins to read, and which to ignore—a challenge that is aggravated by the continuously increasing amount of information that is available online. For instance, many press agencies produce over 12,500 bulletins a year. Reporters therefore have to make selections, and although the agencies often tag their bulletins, the sheer mass of information makes that it is easy to miss important ones.

In this paper, we will focus on one solution to this problem: recommender systems. Typically, corresponding decision aids automatically come up with a pre-selection of information that is worth further consideration, saving institutions, firms, and people parts of the time and effort

otherwise required to separate the relevant from the irrelevant. In particular, here we will evaluate six models that can solve the problem of information selection for the scientific domain.

All models select relevant scientific papers from a large collection of scientific abstracts. They include the *Publication Assistant* (Van Maanen, Van Rijn, Van Grootel, Kemna, Klomp, & Scholtens, in press), a recommender system that was recently developed to assist scientists in identifying relevant articles. We will compare the performance of this system to that of three simple decision heuristics, including a unit-weight linear model (see Dawes, 1979; Dawes & Corrigan, 1974), and two lexicographic rules, called *take-the-best* (Gigerenzer & Goldstein, 1996), and *naïveLex*. We will also pit all models against two complex linear weighted-additive models, one being *Franklin's rule* (Gigerenzer & Goldstein, 1999) and the other multiple regression.

While we do not aim to model the cognitive processes that are actually going on when scientists make literature choices, except for multiple regression all models tested here are grounded in cognitive theories. The *Publication Assistant* is a memory model that is based on the *rational analysis* framework (Anderson, 1990; Oaksford & Chater, 1998), as incorporated in the *ACT-R cognitive architecture* (Anderson & Lebiere, 1998). The heuristics are models of decision making that are grounded in the *fast and frugal heuristics framework* (Gigerenzer, Todd, & the ABC Group, 1999). The linear weighted additive model, Franklin's rule, is also a model of decision making (Gigerenzer & Goldstein, 1999). All models are common in the memory and judgment and/or decision making literature.

In what follows, we will give an outline of the *Publication Assistant*. Next, we will introduce the five alternative models. In an experiment, we will then evaluate the models' performance in predicting scientists' literature preferences.

The Publication Assistant: A Memory Model

An example of the problem addressed in this paper is the selection of relevant talks when attending a large, multi-track scientific conference such as the Annual Cognitive Science Conference. The information selection process starts when a researcher registers and receives a copy of the conference program. For instance, a strategy often employed by many conference attendees is to scan talk titles, author

names, or abstracts for words or names that sound familiar. If an entry contains enough interesting words, it is selected for more careful reading, and the corresponding talk might be attended. In order to determine if a word qualifies as interesting in the context of the conference, a researcher might assess whether she has used the word in her own research in the past. The assumption is that the words used by someone in the context of her own research reflect her scientific interests. The Publication Assistant is a literature selection tool that could be run over a (digitized) conference program prior to attending the conference. The model recommends talks a given scientist might find useful to attend, saving that researcher the time and effort required to scan the conference program on her own. To this end, the model searches through the scientist's own work, examining in how far words that appear in conference abstracts also occur in the scientist's work. Specifically, the model bases its recommendations on the following properties of words in an abstract:

Recency of occurrence in the scientist's own work

- The year in which a word from a conference abstract appears for the first time in the abstracts the scientist has published in the past,
- The year in which a word from a conference abstract appears for the last time in the abstracts the scientist has published in the past,

Frequency of occurrence in the scientist's own work

- The frequency of appearance of a word from a conference abstract in the abstracts the scientist has published in the past,
- The frequency of co-occurrence of a word from the conference abstract with another word in the abstracts the scientist has published in the past.

Based on these properties, the model creates an individual representation of a researcher's interests. The Publication Assistant applies these *user models* to predict the relevance of words that occur in other scientific abstracts, essentially estimating how familiar the contents of these abstracts would be to the scientist. In the next section, we will describe in more detail how the Publication Assistant estimates familiarity.

Model Equations

The Publication Assistant works like a model of the contents of a researcher's memory. Its equations are based on Anderson and Schooler's (1991) rational analysis of memory. According to their analysis, the probability that a *fact* (e.g., a word) stored in memory will be needed to achieve a processing goal can be predicted from the organism's pattern of prior exposure to the corresponding piece of information. For example, the probability that a fact about a scientific topic is of relevance to a researcher may depend on the frequency and recency of his writings about it in the past. Frequency and recency, in turn, feed into a memory currency called *base-level activation*, which

influences a researcher's familiarity with the fact. These relations are captured by Equation 1, in which B stands for the base-level activation of a fact i , t_i stands for the time that has passed since the last exposure to that fact, and d represents the speed with which the influence of each exposure decays away. The summation takes place over all n previous encounters with the fact.

$$B = \ln \left[\sum_{i=1}^n t_i^{-d} \right]. \quad (1)$$

Besides frequency and recency of encounters with facts, the context in which these facts occur also plays a role in the activation of the facts. This *spreading activation* (Quillian, 1968) component represents the likelihood that a fact will be needed if another one is currently being used. These likelihoods depend on the pattern of prior exposures with the facts, as represented by the relatedness measure R_{ji} between two facts j and i (Anderson & Lebiere, 1998):

$$R_{ji} = \frac{F(W_j \& W_i)F(N)}{F(W_j)F(W_i)} \quad (2)$$

where $F(W_j)$ is the frequency of fact i , $F(N)$ is the total number of exposures, and $F(W_j \& W_i)$ is the number of co-occurrences of the facts j and i .

With the equations that are provided by the rational analysis of memory, one can calculate the base-level activation of a word based on its occurrences in publications of the user. However, rather than using Equation 1 directly, the Publication Assistant uses Petrov's (2006) version of it. In Equation 3, the decay parameter is fixed at .5 and a history factor h is added, which represents a free parameter:

$$B = \ln \left[\frac{1}{\sqrt{t_1 + h}} + \frac{2n - 2}{\sqrt{t_n} + \sqrt{t_1 + h}} \right] \text{ with } h > 0 \quad (3)$$

The Publication Assistant makes recommendations by combining the base-level activation of a word (i) with the weighted base-level activation of related words (j) in the abstract (Pirolli & Card, 1999):

$$A_i = B_i + \sum_j B_j R_{ji} \quad (4)$$

To compare the relevance of abstracts with each other, each one is represented by the average activation of the words that occur in it. In a comparison of two abstracts, the Publication Assistant then recommends the more activated one. Abstracts in which many words have high base-level and spreading activation values have a high match with the researchers own word usage, and thus may be more interesting.¹ The Publication Assistant's recommendations are thus based on the structure of the environment of a

¹ Van Maanen et al. (in press) found that the frequency of words in scientific abstracts differs from normal word usage in written English. To counter the unwanted influence of normally high-frequency words (e.g., "the"), van Maanen et al. built a filter for these words when they developed the Publication Assistant. Here, we run all analyses using that filter. As they showed, the filtering does not interfere with how well an abstract represents the contents of a paper.

particular researcher. In particular, the structure of word usage in previously published abstracts. The only parameter that may be varied is the history parameter, h , which represents the relative importance of recency versus frequency in determining activation. In the research reported here, we kept h constant at the same value reported in Van Maanen et al. (in press).

Alternative Models: Decision Strategies

To evaluate the performance of the Publication Assistant in predicting scientists' literature preferences, we compared it to five alternative models. While the Publication Assistant essentially mimics a model of memory, these alternative models have originally been proposed as decision strategies in the judgment and decision making literature

In particular, we focus on a class of models that have been prominent in the fast and frugal heuristics framework. According to this framework, humans often make decisions under the constraints of limited information processing capacity, knowledge, and time—be they about the relevance of scientific articles, or the likely performance of stocks, or the nutritional value of food. Such decisions can nevertheless be made successfully because humans can rely on a large repertoire of simple decision strategies, called heuristics. These rules of thumb can perform well even under the above-mentioned constraints. They do so by exploiting the structure of information in the environment in which a decision maker acts and by building on the ways evolved cognitive capacities work, such as the speed with which the human memory system retrieves information (for recent overviews, see Cokely, Schooler, & Gigerenzer, in press; Marewski, Galesic, Gigerenzer, 2009).

One of the heuristics tested here, the unit-weight linear model, is particularly simple, requiring no free parameters to be fitted. Related models have proved to perform quite well in predicting unknown events and quantities. Just as the unit-weight linear model, also naiveLex dispenses with all free parameters. If these two particularly simple heuristics predicted scientist's literature preferences successfully, then they would simplify the selection of abstracts more than the Publication Assistant does. Take-the-best is a little more complex, requiring free parameters to be fitted for each individual scientist. Take-the-best and related models have been found to be, on average, more accurate than multiple regression in predicting various economic, demographic, and environmental, variables (e.g., Czerlinski, Gigerenzer, & Goldstein, 1999). Finally, the most complex models tested here, Franklin's rule and multiple regression, require for each individual researcher as many free parameters to be fitted as there are words in the abstracts under consideration. While these two models are prominent in the judgment and decision making literature, due to their large complexity they are not considered *heuristic* decision strategies in the fast and frugal heuristics framework. Rather, they are often used as benchmark to evaluate the performance of heuristics in model comparisons (Gigerenzer & Goldstein, 1996).

Lexicographic Heuristics: take-the-best, naiveLex

The first model to be considered here is take-the-best. To make literature recommendations, take-the-best uses attributes of articles as *cues*. In our context, cues are the words that occur in an abstract. If such a word also occurs in a scientist's own publication, then it is considered a *positive* cue, suggesting that an abstract is of interest to that scientist. Take-the-best considers all cues sequentially (i.e., one at a time; hence lexicographic) in the order of their *validity*. The validity of a cue is the probability that an alternative A (e.g., an article) has a higher value on a criterion (e.g., relevance for a researcher) than alternative B, given that alternative A has a positive value on that cue and alternative B does not. In a comparison of two abstracts, take-the-best bases a decision on the first cue that *discriminates* between the abstracts, that is, on the first cue for which one abstract has a positive value and the other one does not. The heuristic is defined in terms of three rules:

- (1) Look up cues in the order of their validity.
- (2) Stop when the first cue is found that discriminates between the abstracts.
- (3) Choose the abstract that this cue favors.

The second lexicographic model, here called naiveLex, is identical to take-the-best, except that it does not estimate the validity order of cues. Rather, cues are simply considered in the order of the frequency of occurrence of the corresponding words in each researcher's published abstracts. This aspect of the model is similar to the Publication Assistant, in which the word frequency is also taken into account (but weighted with recency).

A Unit-weight-linear Heuristic

Lexicographic heuristics such as take-the-best can avoid going through all words (i.e., cues) from an abstract, which can save effort, time, and computations once the order of cues is known. Unit-weight linear heuristics, in contrast, integrate all cues into a judgment by adding them. These models can nevertheless simplify the task by weighing each cue equally (hence unit-weight). In a comparison of two abstracts, it reads as follows:

- (1) For each abstract, compute the sum of positive cues.
- (2) Decide for the abstract that is favored by a larger sum.

Weighted-Additive models: Franklin's Rule and Multiple Regression

Franklin's rule (Gigerenzer & Goldstein, 1999) is similar to the unit-weight linear heuristic, but instead weights all the cues by their validities prior to summation. (The cue validities are identical to those relied on by take-the-best.) Multiple regression, in turn, estimates the weights of the cues by minimizing the error in the calibration set using maximum likelihood estimation. In a comparison of two abstracts, Franklin's rule and multiple regression read as:

- (1) For each abstract, compute the weighted sum of cues.
- (2) Decide for the abstract that is favored by a larger sum.

Experiment

To compare the Publication Assistant to the alternative models' capability of predicting actual scientist's literature preferences, we re-analyzed data from a study by Van Maanen et al. (in press, Experiment 2). They had asked researchers from the field of cognitive science to rate how much they were interested in a paper after reading the abstract. In this study, Van Maanen et al. had found that the Publication Assistant could fit researcher's interests reasonably well; however, they did not compare its performance to that of alternative models.

Methods

Participants Ten researchers (2 full professors, 2 associate professors, 5 assistant professors, and 1 post-doc) from various subfields of cognitive science and from various countries were asked to participate.

Procedure For each of the researchers, Van Maanen et al. (in press) constructed user models of the Publication Assistant based on the abstracts of their published work insofar it was indexed by PsycINFO. They then ordered all abstracts from the last three volumes (2004-2006) of the *Cognitive Science Journal* according to the predicted relevance for the researcher, based on the researcher's published abstracts. From the ordered list of abstracts, they presented each researcher the top five abstracts, the bottom five abstracts, and five abstracts from the middle of the list. For each researcher, the presentation order of these 15 abstracts was randomized. Each researcher indicated with a grade between 0 and 9 how much he or she was interested in the papers, based on the abstracts.

Analyses When comparing models that differ in terms of their complexity, it is advisable to assess the models' ability to generalize to new data (e.g., Marewski & Olsson, 2009; Pitt, Myung, & Zhang, 2002). This holds especially true when the models have been *designed* to generalize to new data, as it is the case with the recommender systems tested here. To compare the performance of the Publication Assistant to that of the five alternative models in predicting each researcher's ratings, we ran a cross-validation. To this end, we constructed paired comparisons of all 15 abstracts for each participant individually (210 pairs). We divided each participant's abstracts pairs randomly into two parts. The first part represented the *calibration set* in which we calculated for each participant that person's optimal values for the free parameters in take-the-best, Franklin's rule, and multiple regression, respectively. That is, we identified the parameter value at which each model would correctly




Figure 1. The performance of the non-calibrated models. A-J represent individual participants. PA: Publication Assistant; UWL: unit-weight linear heuristic.

predict the largest proportion of literature preferences. Take-the-best, Franklin's rule, and multiple regression will therefore be referred to as the calibrated models.

We used these optimal values to compute the proportion of preferences consistent with each model in the other half, the *validation set*, where the models' generalizability is evaluated. For each partition, we also computed the proportion of preferences consistent with the three non-calibrated models (the Publication Assistant, naiveLex, and the unit-weight linear heuristic). The free parameter of the Publication Assistant, h , was set to 10. In fitting the very same participants as we do here, van Maanen et al. (in press), had found this value to work reasonably well.

We ran these analyses for a subset of possible sizes of the calibration and validation sets; that is, we first computed the proportion of each model's correct predictions for a calibration set size of 1 and a test set size of 209, then for a calibration set size of 11, and a test set size of 199, and so on. The larger the size of the calibration sets, the larger is the sample of paired comparisons from which the parameterized decision models can estimate an individual researcher's interests, that is, the more "experience" these models can accumulate before making their predictions. This procedure was repeated enough times to average out noise due to the random selection of calibration sets.

Results When comparing the Publication Assistant with the other non-calibrated models (naiveLex and the unit-weight linear model), we found that the three models performed differently for different participants (Figure 1). The Publication Assistant made the most correct inferences for three participants (A, D, and J), while unit-weight linear heuristic outperformed the other two non-calibrated competitors on four occasions (B, C, E, and H). NaiveLex scored best for three participants (F, G, and I). Overall, the performance of the models did not differ much ($M_{PA} = .60$, $M_{naiveLex} = .59$, $M_{UWL} = .60$).

For each of the 10 participants, Figure 2 shows the proportion of correctly predicted preferences for the three calibrated models as a function of the size of the calibration set. As one would expect, for all participants the accuracy of the predictions of the parameterized models increases with the size of the calibration set. Of the calibrated models, Franklin's rule was consequently outperformed by the take-the-best heuristic and the multiple regression model, which

performed equally well, but differed among participants. Take-the-best was the best predictor for participants B, C, E, G, I, and J, while the regression model performed best for participants A, D, F, and H. Overall, take-the-best performed best ($M_{TTB} = .84$, $M_{MR} = .81$, $M_{Franklin} = .71$).

Discussion

We examined the ability of six models to predict scientists' literature preferences: (i) the Publication Assistant, a recommender system that is based on a rational analysis of memory and the ACT-R architecture; (ii-iv) three simple heuristics, including take-the-best, a naive lexicographic model, and a unit-weight linear model, and (v-vi) two complex weighted-additive models, Franklin's rule and multiple regression.

For some participants and calibration set sizes, the regression model outperformed take-the-best. One reason why take-the-best did not fare as well as multiple regression on every occasion might be that the structure of information available in the abstracts was not well suited for this simple heuristic (see Martignon & Hoffrage, 2002). For instance, take-the-best essentially bets on a noncompensatory information structure, always preferring the most valid discriminating cue to all others. In the domain of literature selection, such information structures might not be prevalent. To give an example, the words "Memory" and "Retrieval" might be equally good predictors of some cognitive scientist's research interests.

Another result was that the performance of the non-

calibrated models varied across participants. However, it should be realized that naiveLex and the Publication Assistant only differ with respect to the use of the recency component. Both models use the frequency of words in published abstracts in the same way. Therefore, the variation in performance between the models may be attributable to the importance the recency component plays in the models. That is, the Publication Assistant overestimates the importance of more frequent words in the published abstracts relative to the importance of recent word usage. Since the h factor scales the relative contributions of word frequency and word recency, recommendations of the Publication Assistant could improve if one would allow for the h parameter to be fit individually to each scientist's data.

The fact that the non-calibrated models performed differently across participants is in agreement with other findings in the judgment and decision making literature, where individual variation in people's use of decision strategies is commonly observed (Bröder & Gaissmaier, 2007; Mata, Schooler, & Rieskamp, 2007; Pachur, Bröder, & Marewski, 2008). In fact, based on this literature, one might expect that the most useful approach for designing recommender systems would have been to build different systems for different users, depending on which model predicts the respective scientist's preferences best.

Our results also complement findings by Lee, Loughlin, and Lundberg (2002), who, in a study on literature search, examined the performance of a simple heuristic in identifying articles that are relevant to a given topic (e.g.,




Figure 2. The calibrated models' individual predictions of literature selections. Each panel represents one participant. TTB: take-the-best; Franklin: Franklin's rule; MR: multiple regression.

memory). Their analyses show that a researcher going by a variant of take-the-best would have had to search through fewer articles in order to find the relevant ones than a person behaving in accordance with a weighted-additive model.

Why was the Publication Assistant outperformed?

Take-the-best, Franklin's rule, and the regression model learned about the scientists' interests directly from the paired comparisons between abstracts that were included in the calibration sets. The Publication Assistant, in turn, was trained on a participant's published abstracts (Van Maanen et al. in press), under the assumption that word frequencies in those abstract would reflect the participants' interests. While this way of training the model better reflects real-life situations of information selection, in which people's appraisal of items (such as abstracts) is often unknown, it might have been detrimental for the model's performance.

Conclusion

In this paper, we evaluated the ability of models of memory and decision making to serve as literature recommender systems. As it turned out, the performance of the models in predicting literature preferences differed substantially across participants, indicating that there may be substantial individual differences in the ways scientists decide which papers to read. This finding suggests that for successful recommendation, the best predicting model should be determined first on an individual basis.

To conclude, in today's world of mass media, the choice which information to attend to, and which to ignore becomes an ever more important challenge for professionals. Automatic recommender system might help to cope with these demands of the information age—savings in time and effort that can eventually be invested elsewhere. We hope that comparisons between different approaches, such as the ones tested here, help along that way.

References

- Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Erlbaum.
- Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Lawrence Erlbaum.
- Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science*, 2(6), 396-408.
- Bröder, A., & Gaissmaier, W. (2007). Sequential processing of cues in memory-based multiattribute decisions. *Psychonomic Bulletin & Review*, 14(5), 895-900.
- Cokely, E.T., Schooler, L.J., & Gigerenzer, G. (in press). Information use for decision making. In M.N. Maack & M.J. Bates (Eds.), *Encyclopedia of Library and Information Sciences*.
- Czerlinski, J., Gigerenzer, G., & Goldstein, D. G. (1999). How good are simple heuristics? In G. Gigerenzer, P. M. Todd & the ABC Group (Eds.), *Simple heuristics that make us smart* (pp. 97-118). New York: Oxford UP.
- Dawes, R. M. (1979). Robust beauty of improper linear-models in decision-making. *American Psychologist*, 34(7), 571-582.
- Dawes, R. M., & Corrigan, B. (1974). Linear-models in decision-making. *Psychological Bulletin*, 81(2), 95-106.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4), 650-669.
- Gigerenzer, G., & Goldstein, D. G. (1999). Betting on one good reason: The take the best heuristic. In G. Gigerenzer, P. M. Todd & the ABC Group (Eds.), *Simple heuristics that make us smart*. New York: Oxford UP.
- Gigerenzer, G., Todd, P. M., & the ABC Group. (1999). *Simple heuristics that make us smart*. New York: Oxford UP.
- Lee, M. D., Loughlin, N., & Lundberg, I. B. (2002). Applying one reason decision-making: The prioritisation of literature searches. *Australian Journal of Psychology*, 54(3), 137-143.
- Martignon, L., & Hoffrage, U. (2002). Fast, frugal, and fit: Simple heuristics for paired comparison. *Theory and Decision*, 52(1), 29-71.
- Marewski, J. N., Galesic, M., & Gigerenzer, G. (2009). Fast and frugal media choices. In T. Hartmann (Ed.), *Media choice: A theoretical and empirical overview* (pp. 107-128). New York & London: Routledge.
- Marewski, J. N., & Olsson, H. (2009). Beyond the null ritual: Formal modeling of psychological processes. *Journal of Psychology*, 217, 49-60.
- Mata, R., Schooler, L. J., & Rieskamp, J. (2007). The aging decision maker: Cognitive aging and the adaptive selection of strategies. *Psychology and Aging*, 22, 796-810.
- Oaksford, M., & Chater, N. (1998). *Rational models of cognition*. Oxford: Oxford UP.
- Pachur, T., Bröder, A., & Marewski, J. N. (2008). The recognition heuristic in memory-based inference: Is recognition a non-compensatory cue? *Journal of Behavioral Decision Making*, 21(2), 183-210.
- Petrov, A. A. (2006). Computationally efficient approximation of the base-level learning equation in ACT-R. In D. Fum, F. Del Missier & A. Stocco (Eds.), *Proceedings of the Seventh International Conference on Cognitive Modeling* (pp. 391-392). Trieste, ITA.
- Pirolli, P., & Card, S. (1999). Information foraging. *Psychological Review*, 106(4), 643-675.
- Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method for selecting among computational models for cognition. *Psychological Review*, 109, 472-491.
- Quillian, M. R. (1968). Semantic memory. In M. Minsky (Ed.), *Semantic information processing* (pp. 216-270). Cambridge, MA: MIT Press.
- Van Maanen, L., Van Rijn, H., Van Grootel, M., Kemna, S., Klomp, M., & Scholtens, E. (in press). Personal Publication Assistant: Abstract recommendation by a cognitive model. *Cognitive Systems Research*.