A Comparison of Rule-Based versus Exemplar-Based Categorization Using the ACT-R Architecture

Matthew Rutledge-Taylor, Christian Lebiere, Robert Thomson, James Staszewski, and John R. Anderson
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15218
412-268-6025
mfrt@andrew.cmu.edu, cl@andrew.cmu.edu, thomsonr@andrew.cmu.edu, jjs@cmu.edu, ja@cmu.edu

Keywords:
Modeling, Categorization, Learning

ABSTRACT: *A rule-based approach to categorization is compared with an exemplar-based approach. Both models were developed using the ACT-R architecture. Both approaches yield similar accuracy and are relatively impervious to varying model parameters. Implications for the nature of implicit and explicit knowledge and learning are discussed.*

1. Overview

The research in this paper pertains to the general issue of the relationship between existing cognitive theories of categorization. The use of cognitive models that implement descriptive theories allows for quantitative comparisons of the theories. ACT-R has been used previously to compare and contrast exemplar-based and rule-based approaches to categorization (Anderson & Betz, 2001). This paper describes two groups of ACT-R models of a common task. One set is exemplar-based while the other is rule-based. The results of this paper support the notion that exemplar-based theories and rule-based theories are largely commensurable.

Two groups of ACT-R models of a categorization task are compared. The task was to categorize four types of facilities depicted in simulated satellite images. Each facility type corresponded to one category; the terms facility type and facility category will be used interchangeably, hereafter. Each facility type was defined in terms of the probabilities of the presence or absence of various facility features, none of which were unique to any facility. Human participants were trained to learn the facility categories by studying multiple examples of each. One group of ACT-R models learned the categories by storing the (declaratively represented) examples in memory (i.e., exemplar models). The models in the other group were provided with explicit rules that were supplied to the models a priori.

Within each group of models, there were versions of the model that attended to different features in the examples; in the case of the exemplar models, these versions corresponded to different ACT-R memory retrieval mechanisms (partial matching and spreading activation). The purpose of this modeling effort was to compare the performance of rule-based ACT-R models of categorization to exemplar-based ACT-R models of the task described. The fact that this effort was successful has interesting implications beyond this specific project as it provides evidence for the general commensurability of exemplar and rule-based theories of categorization.

1.1 Category Learning

There are many distinct theories of category learning. Most fall into three main groups: rule-based theories (Goodman, Tenenbaum, Feldman & Griffiths, 2008), prototype theories (Rosch, 1973), and exemplar theories (Nosofsky, 1986).

Rule-based theories are committed to the ability of categorizers to identify the category of an object (or an abstract concept) by testing it against one or more rules. Rules typically take an if/then form whereby the object is deemed to be a member of a category (or, is ruled out) if it satisfies the 'if' conditions of one or more rules. Rule-based theories, such as RULEX, can include the possibility of exceptions (Nosofsky & Palmeri, 1995; Nosofsky, Palmeri & McKinley, 1994; Palmeri & Nosofsky, 1998) and/or probabilistic category assignment (Goodman et al., 2008). The rule-based ACT-R model discussed below employs rules to determine the probabilities that an unlabeled facility is a member of each of four possible categories.

Prototype theory postulates that learned categories are represented, mentally, by prototypes. The membership of an instance to a category is determined by the agreement between the properties of the prototype and the properties of the instance. Multiple-prototype theories allow for multiple prototypes for each category to accommodate non-linearly separable categories.

Exemplar theories postulate that category instances (i.e., exemplars) are memorized individually. Category assignment decisions are made by comparing a new instance to existing exemplars. For example, the judged category can be the one belonging to the exemplar nearest (most similar) to the new instance (i.e., winner-take-all); alternately, the judged category can be based on a function of the combined distances from the new instance to each of the exemplars (e.g., least mean squared distance for each set of category exemplars). Standard declarative memory retrieval in ACT-R would support a winner-take-all categorization process. However, the exemplar-based ACT-R model discussed below makes use of a mechanism called blending, which allows all exemplars to contribute to categorization decisions.

1.2 ACT-R Architecture

ACT-R is a computational implementation of a unified theory of cognition (Anderson et al., 2004; Anderson & Lebiere, 1998). It accounts for information processing in the mind via a set of task-invariant mechanisms, which are constrained by biological limitations of the brain. It consists, primarily, of a set of modules, such as the declarative memory system (DM), and a production system. Each module exposes a buffer, which contains a single chunk, to the rest of the system. Each chunk is a member of a chunk type, and consists of a set of type-defined slots with instance specific values.

Information is processed in ACT-R by the production system, which operates on the contents of the buffers. Each production consists of an if-then condition-action pair. Conditions are typically criteria for buffer matches, while the actions are typically changes to the contents of buffers or actions that trigger operations in the associated modules (e.g., the recalling of a memory). The normal process sequence for a model is to loop through probabilistically selecting an eligible production to fire and executing its effects on the system until no production matches the state of the system, causing the model to stop. The production with the highest net utility (after the effects of noise are factored in) is selected to fire from among the eligible productions.

When a retrieval request is made to declarative memory, the single most active matching chunk is returned. Chunk activation is computed as the sum of base-level activation, spreading activation, mismatch penalty and stochastic noise (see figure 1). Spreading activation is a mechanism that propagates activation from the contents of buffers to declarative memory proportionally to their strength of association. The consequence of this is that chunks in DM that share content with chunks in buffers will have an increased probability of being recalled irrespective of degree of match. Partial matching is a mechanism that allows for chunks in memory that do not perfectly match a retrieval request to be recalled if their activation overcomes a similarity-based mismatch penalty.

$$A_i = B_i + S_i + P_i + \mathcal{E}_i$$

Figure 1. The chunk activation formula in ACT-R. A_i is the net activation, B_i is the base-level activation, S_i is the effect of spreading activation, P_i is the effect of the mismatch penalty, and $\mathcal{E}_i$ is magnitude of stochastic noise.

An advanced memory retrieval mechanism, called blending, differs from standard retrieval in that all chunks in DM that match the retrieval request specification are blended together to create a new chunk, which is retrieved (Lebiere, 1999). This mechanism allows for exemplar categorization models similar to those described in Shi et al. (2010), to be created in ACT-R. These models obey the Luce choice axiom (see figure 2; Luce, 1959), where the weight of each exemplar is based on a similarity metric. The default similarity metric is to compare chunk slot values. In the case of the models discussed in this paper, this amounts to comparing the occurrences of facility features (discussed below).

$$P(i) = \frac{w_i}{\sum_j w_j}$$

Figure 2. Luce choice axiom. The probability that option i is selected is relative to the weighted sum of the pool of options j.

2 Facility Identification Task

Experimental participants were trained to identify four kinds of facilities in simulated geospatial images. Each image is of a single facility (e.g., factory complex) that is composed of a collection of discrete features (e.g., buildings) drawn, probabilistically, from three distinct categories. The three categories of features were: IMINT (image intelligence), representing buildings and other terrain features such as roads and rivers; MASINT (measurement and signature intelligence), representing signals such as chemical concentrations or radiation etc.; and, SIGINT (signals intelligence), representing communication transmissions. There were nine unique IMINT features, seven that represented buildings, and two that represented water features. In contrast, there were only two kinds of MASINT features, while the SIGINT features were entirely homogeneous. Each IMINT could appear at most one time in each image, whereas multiple instances of SIGINT and each MASINT could occur in each image. Additionally, each building (IMINT) could have one piece of rooftop hardware attached to it, or not.

The four facilities were defined by different probabilities for the occurrences of each of the possible features. In the case of the IMINT features, these probabilities simple defined the likelihood of the feature occurring in an instance of the given facility. In the case of MASINTs, SIGINTs, and rooftop hardware the probabilities defined the likelihoods of few or many instances of the feature.

The experiment was divided into two main phases: a training phase and a testing phase. In the training phase the participants were presented with 48 annotated examples of each facility (192 total examples), 16 at a time (in a four by four grid). Participants were not limited in how long they could study the images. Training time ranged from 8 minutes to 73 minutes (mean 24 minutes). In the testing phase the participants were presented with single unlabeled images, one at a time. For each image, the participant was required to report a probability distribution over the four possible facilities indicating the likelihood that the image contained each of the facilities.

3 ACT-R Models

This paper is devoted to comparing and contrasting the performance of ACT-R models of the facility identification task. The main comparison is between models that instantiate an exemplar theory account of category learning and models that instantiate a rule-based account. In common to all the models discussed below are the following details.

The testing phase in the simulations consisted of the presentation and categorization of 300 simulated images of unknown facilities. For each presentation, the ACT-R model holds an instance of a facility frame representing the current facility under examination in the imaginal buffer, which corresponds to the parietal lobe of the brain (Anderson, 2007). The facility chunk-type defines a slot for the facility type, a slot for the total number of IMINTs in the image, a slot for the total SIGINTs, one slot each for the totals of two kinds of MASINT, a slot for the total number of pieces of rooftop hardware on buildings in the image, and one slot for each of the nine kinds of IMINT. Each IMINT slot stores a chunk representing the presence of that IMINT or is left empty. There were an average of 5.05 IMINT features per facility instance.

Three related concepts are used almost interchangeably in this paper: facility chunk, facility frame, and facility exemplar. A facility chunk is an ACT-R representation of a facility instance. Facility exemplars are represented as chunks in declarative memory. The term facility frame is used to refer to schematic structure of a facility and is used in the context of maintaining a facility in an ACT-R buffer.

3.1 Exemplar Models

During the training phase the annotated images are imported into the declarative memory of the model one at a time. For each image, the model temporarily holds a facility frame in working memory by populating the imaginal buffer with appropriate chunk representations of the features present in the image. Once filled, the imaginal buffer is cleared and the facility chunk is committed to memory (DM).

During the testing phase, images are presented to the model one at a time. The model performs a blended retrieval request of DM for a facility frame chunk based on some information available in the images. The facility slot value of the blended chunk is used as the model's answer to the identification question. The model is able to assign probabilities to each category by converting the activations of all exemplar chunks of each category to an aggregate probability using the ACT-R Probability Retrieval (Boltzmann) Equation (see figure 3). Note that the Boltzman equation below is equivalent to the Luce choice axiom (shown above in figure 2), where w = $e^{A/s}$.

$$P_i = \frac{e^{A_i/s}}{\sum_j e^{A_j/s}}$$

Figure 3. Boltzman equation controlling the probability (P_i) that chunk i with activation A_i is retrieved relative to the activation levels of all of eligible chunks (j). The s parameter reflects chunk activation noise.

Two distinct retrieval mechanisms could apply to the recall of facility frames. They are partial matching and spreading activation.

3.1.1 Partial Matching Model

The partial matching version of the ACT-R model uses only the slots that store the counts of the various feature types (and hardware) as part of the retrieval request. This model represents a participant who is not attentive to the particular buildings that are present in a test image. When classifying a facility image, the model compares the feature counts in the image to the counts in facility chunks in DM (see figure 4).

The effect of partial matching is that the model is able to make similarity-based inferences in making facility discriminations. By limiting the model to representing only the numbers of each feature type, this similarity-based inference mechanism is fruitful only if there is a statistically significant difference in the distribution of feature totals for the different facilities. Our results show that such a statistical relationship exists.

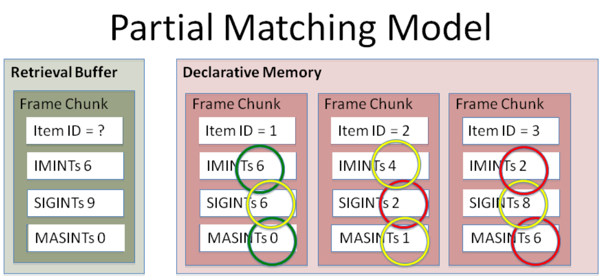

Figure 4. Green circles indicated perfect matching slots values; yellow circles indicated good matches; and, red circles indicate poor matches.

The ACT-R model making use of partial matching only was able to correctly identify the facility in each test sector 46.2% of the time, on a cross-validated 80%/20% training/testing split of the 300 sample scenes. The confusion matrix (see Table 1) listing the probability of classifying an instance of a given facility type as any of the four facility type options shows a pattern dominated by confusion between facilities A and C, and B and D.

Table 1
Confusion matrix for partial matching model

Facility	*A*	*B*	*C*	*D*
A	**.559**	.090	.274	.077
B	.077	**.490**	.116	.316
C	.356	.124	**.375**	.145
D	.108	.288	.180	**.424**

3.1.2 Spreading Activation Model

The spreading activation version of the model ignores the feature counts; instead, the IMINT features in the imaginal buffer form the context of retrieval. The model assembles a facility frame in the imaginal buffer using chunks representing the IMINT features present in the image. Each feature spreads activation to facility chunks in DM that include the feature to a degree inversely proportional to the logarithm of the number of chunks including that feature (see figure 5), a phenomenon known as the fan effect (Anderson, 1974; Rutledge-Taylor & West, 2008; West et al. 2010). This model represents a participant who is solely focused on the particular buildings in the image. When the request for a facility chunk from DM is made, chunks that share IMINT features in common with the image will get a boost in activation, increasing the probability that they will be retrieved.

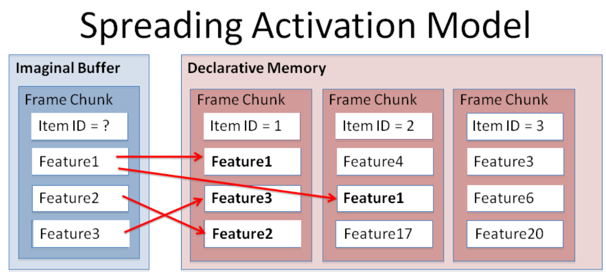

Figure 5. Frames that share features in common with the context in the imaginal buffer receive a boost in activation.

Performance of the spreading activation version of the model was better than the partial matching version. It was able to correctly identify the facility 65.5% of the time, on a cross-validated 80%/20% training/testing split of the data. The pattern of individual confusion probabilities (see Table 2) reflects the overlap between the features likely to belong to each facility, as well as the fact that the number of features increases from facility A to B, C and D, leading to more spreading activation for the latter.

Table 2
Confusion matrix for spreading activation model

Facility	*A*	*B*	*C*	*D*
A	**.585**	.017	.247	.151
B	.006	**.635**	.061	.297
C	.065	.062	**.585**	.287
D	.025	.108	.054	**.813**

3.1.3 Combined Model

A third version of the model, referred to as the combined model, which uses both partial matching over feature counts and spreading activation from IMINT features, was created. In this model, frames include slots for both the feature totals, and for each IMINT feature individually. The count slots are used as retrieval cues (partial matching), while the IMINT features are used as retrieval context (spreading activation). Performance on this was somewhat better than for either model alone. It was able to correctly identify the facility 72.0% of the time, on a cross-validated 80%/20% training/testing split of the data. The confusion matrix is presented in table 3 and displays uniformly good performance.

Table 3
Confusion matrix for combined model

Facility	*A*	*B*	*C*	*D*
A	**.719**	.013	.177	.090
B	.006	**.716**	.074	.203
C	.065	.058	**.691**	.185
D	.035	.133	.079	**.753**

3.1.4 Blending

By default in ACT-R, a retrieval request to declarative memory produces the single chunk representing the frame with the greatest net activation. An exemplar-based model using standard retrieval would be a winner-take-all categorization model. However, we hypothesize that the experimental participant is not making the facility identification judgment based on the single exemplar in memory that best matches the set of features in the target sector. Rather, every exemplar in memory should contribute to the categorization decision. The relative contribution of a chunk is a function of its base-level activation, partial matching and spreading activation. The blending mechanism creates a new chunk of the requested chunk-type that is an aggregate of all the exemplars in memory. The value for each of the new chunk's slots is that which is the best compromise value amongst all the values occurring in all the exemplars, weighted by the activation strength of the exemplars. The blended facility category slot value corresponds to the model's categorization decision.

The outcome of blending is somewhat different than generating a prototype in that the relative contribution of each exemplar is based on its activation, which is affected by the specific retrieval cues (i.e., the specific features present in the facility to be identified). Additionally, there are no persistent prototypes in DM. This is why we consider these models to implement exemplar-based categorization. However, a case can be made that this could be considered a kind of dynamic multiple-prototype learning, if the persistence of prototypes is not necessary.

3.2 Rule-based Models

The ACT-R models of rule-based category learning presented in this paper were created as a proof of concept that a particular choice of rule representation would be effective in producing categorization accuracy approximately equal to, or better than, the exemplar-based ACT-R models. As such, the models do not learn the categorization rules. Rather, optimal rules were assigned to the models. Each rule specified a layer type, the facility to which the rule applied, a multiplicative likelihood factor, and either a single IMINT feature or in the case of rules about countable features (e.g., SIGINTs), a matching quantity. The condition for the rule match is either the presence of a single IMINT feature, or a quantity of countable INTs. The multiplicative factors for the IMINT features were based on a statistical information gain measure for each feature. The count rule factors were estimated, and parameterizable. An additional parameter of the model was the degree of permissible mismatch between the number of count features specified in a rule and the count in the facility to be identified. For example, the rule chunk, (s1 isa rule layer sigint category a value 4.8 factor 3), encodes the rule that the posterior probability of the unknown facility being of category A is three times greater than the prior probability if the facility is within a threshold difference of 4.8 SIGINT features. The threshold and factors were manipulated experimentally. However, a broad range of thresholds and factors results in near ceiling performance in model accuracy.

The models maintained probabilities for each facility category in the goal buffer, and adjusted these probabilities according to the multiplicative factors encoded in all the rules matching the contents of the sector under examination. For each sector to identify, the model applied all of the applicable rules to produce a final probability distribution. The facility assigned the highest probability was interpreted as the model's forced choice response for accuracy evaluation.

4 Results and Discussion

The performance profiles of the rule-following models were largely parameter-invariant. As such, the results for the various combinations of parameters will not be reported; rather, the results presented below reflect the single set of parameters that best matched the accuracy results of the exemplar-based ACT-R model. The count mismatch threshold was 30% (over or under the value specified in the rule), the IMINT rule multiplicative factor was 1.2, the SIGINT factor was 3.0, the MASINT factor was 3.0, and the hardware count factor was 1.2. Table 4 summarizes the relative accuracies of the exemplar-based and rule-based ACT-R models.

Table 4
Comparison of Rule-Based and Exemplar Models

	Rule-Based	*Exemplar*
PM	.476	.462
SA	.657	.655
Both	.755	.720

The version of the rule-following model analogous to the partial matching version of the statistical learning model only applied rules that pertained to object counts. It scored a 47.6% accuracy rate in categorizing unseen sectors, compared to the 46.2% rate of the exemplar model. Another version of the rule-following model analogous to the spreading activation model only applied rules about the presence of specific IMINT features. It scored an accuracy rate of 65.7% compared to 65.5% for the exemplar model. The rule-following model that applied all rules scored 75.5%, while the exemplar model that attended to counts as well as specific IMINTs scored 72.0%.

The agreement between the exemplar learning ACT-R model and the rule-following ACT-R model support the hypothesis that the rules created for the rule following model captured the same information learned by the exemplar learning model. We hypothesize that this is the case because both the exemplar models and rule-based models maximally exploit the information that can be extracted from the data given the parallel limitations imposed of the models.

Human experimental participants scored a mean accuracy of 53.5%. In post-experiment interviews it was revealed that some participants were explicitly aware of the relationships between feature counts and facility categories. However, Given that they outperformed the partial matching models, it is likely that they were also able to detect correlations between specific IMINTs and facility categories. These correlations would have to be applied in an incomplete or imperfect manner as the human participants scored less than the predicted 72% (or better) accuracy of the combined model. An intriguing possibility is that the participants employed a strategy of augmenting an exemplar-based IMINT feature representation with the application of explicit feature count rules.

Experimental evidence suggests that exemplar theories and rule-based theories can make similar predictions, with each accounting for different phases of concept learning (Rouder & Ratcliff, 2006). Exemplars are be relied upon initially; rules are then inferred from the data; and finally, some exemplars are retained to account for rule exceptions. See Anderson & Betz (2001) for an account of how exemplar-based and rule-based models can be combined to produce classification behavior.

The performance equivalence between the two groups of models establishes that functional Bayesian inference can be accomplished in ACT-R either through explicit, rule application or through the implicit, subsymbolic processes of the activation calculus, that support the exemplar model. This should be expected as the semantics and learning mechanisms of the subsymbolic system in ACT-R is fundamentally Bayesian in nature (Anderson, 1990; 1993). Elsewhere it is argued that exemplar models, interpreted as performing importance sampling, provide a plausible mechanism for implementing Bayesian inference (Shi, et al., 2010). This further supports the notion that the blending mechanism of ACT-R, necessary for the current exemplar model, is a cognitively sound alternative to standard memory retrieval in ACT-R.

The comparison of the sets of functionally equivalent models described in this paper is significant as it provides a principled quantitative comparison between two theoretically distinct accounts of categorization. When implemented in a cognitive architecture that obeys a variety of meaningful constraints the two theories, exemplar-based and rule-based, produced equivalent results. We do not interpret this outcome as supporting one theory over another. Rather, we take it to be a lesson when building models that incorporate a categorization component. Specifically, we let the task dictate whether a rule-based or exemplar-based account is most appropriate, rather than a preconceived notion about how categorization ought to be done cognitively.

5 Acknowledgement

This work is supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of the Interior (DOI) contract number D10PC20021. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained hereon are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI, or the U.S. Government.

6 References

Anderson, J. R. (1974). Retrieval of prepositional information from long-term memory. *Cognitive Psychology, 6*, 451-474.

Anderson, J. R. (1990). *The adaptive character of thought*. Erlbaum, Hillside, NJ.

Anderson, J. R. (1993). *Rules of the mind*. Erlbaum, Hillside, NJ.

Anderson, J. R. (2007) *How Can the Human Mind Occur in the Physical Universe?* New York: Oxford University Press.

Anderson, J.R., & Betz, J. (2001). A hybrid model of categorization. *Psychonomic Bulletin & Review, 8*, 629-647

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., Qin, Y. (2004) An integrated theory of Mind. *Psychological Review, 111*, 1036-1060.

Anderson, J. R., and Lebiere, C. (1998), *The atomic components of thought*, Erlbaum, Mahwah, NJ, 1998

Goodman, N. D., Tenenbaum, J. B., Feldman, J. & Griffiths, T. L., (2008). A Rational Analysis of Rule-Based Concept Learning. *Cognitive Science,* 32, 108-154.

Lebiere, C. (1999). The dynamics of cognitive arithmetic. *Kognitionswissenschaft* [Journal of the German Cognitive Science Society] Special issue on cognitive modelling and cognitive architectures, D. Wallach & H. A. Simon (eds.)., 8 (1), 5-19.Luce, R. D. (1959). *Individual Choice Behavior: A Theoretical Analysis*. Wiley, New York.

Nosofsky, R. M. (1986). Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General*, 115, 39-57.

Nosofsky, R. M. & Palmeri, T. J. (1998). A rule-plus-exception model for classifying objects in continuous-dimension spaces. *Psychonomic Bulletin & Review 5*(3), 345-369.

Nosofsky, R. M., Palmeri, T. J., and McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, 101(1): 53–79.

Palmeri, T. J. & Nosofsky, R. M. (1995). Recognition memory for exceptions to the category rule. *Journal of Experimental Psychology: Learning, Memory & Cognition, 21*(3), 548-568.

Rouder, J. N., & Ratcliff, R. (2006). Comparing exemplar- and rule-based theories of categorization. *Current Directions in Psychological Science*, *15*, 9-13.

Rosch, E. H. (1973). Natural categories. *Cognitive Psychology*, 4, 328-350.

Rutledge-Taylor, M. F. & West, R. L. (2008) Modeling The fan effect using dynamically structured holographic memory. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society*. 385-390. Washington, DC: Cognitive Science Society.

Shi, L., Griffiths, T. L., Feldman, N. H., and Sanborn, A. N. (2010). Exemplar models as a mechanism for performing Bayesian inference. *Psychonomic Bulletin & Review*, 17(4), 443-464.

West, R. L., Pyke, A. A., Rutledge-Taylor, M. F. & Lang, H. (2010) Interference and ACT-R: New evidence from the fan effect. *In Proceedings of the Tenth International Conference on Cognitive Modeling*. Philadelphia, PA: Drexel University.

Author Biographies

MATTHEW RUTLEDGE-TAYLOR is a Post-Doctoral Researcher at Carnegie Mellon University in Pittsburgh, PA. His research interests include modeling theory as well as model building using ACT-R and Holographic Reduced Representations.

CHRISTIAN LEBIERE is a Research Psychologist at Carnegie Mellon University in Pittsburgh, PA. His main research interests are computational cognitive architectures and their applications to psychology, artificial intelligence, human-computer interaction, decision-making, intelligent agents, cognitive robotics and neuromorphic engineering.

ROBERT THOMSON is a Post-Doctoral Researcher at Carnegie Mellon University in Pittsburgh, PA. He received his PhD in Cognitive Science from Carleton University (Ottawa, Canada) in 2011.

JAMES STASZEWSKI is a Professor of Psychology at Carnegie Mellon University in Pittsburgh, PA.

JOHN R. ANDERSON is a Professor of Psychology at Carnegie Mellon University in Pittsburgh, PA.