Experience and Pseudo-Experience: Exemplar Effects Without Feedback

Henrik Olsson (henrik.olsson@psy.umu.se)
Department of Psychology, Umeå University
SE-901 87, Umeå, Sweden

Peter Juslin (peter.juslin@psy.umu.se)
Department of Psychology, Umeå University
SE-901 87, Umeå, Sweden

Abstract

Many real world situations do not offer unambiguous outcome feedback on how to categorize objects. Models in the categorization literature have mostly been formulated for tasks with trial-by-trial outcome feedback. We examined if there was evidence for exemplar memory also when no external feedback is provided and the criterion is derivative of more abstract knowledge. In a "teacher-student" task, a teacher learns how to judge the toxicity of bugs from external outcome feedback and conveys this knowledge to a student that receives no outcome feedback. The results showed that the students exhibit exemplar effects even if the instructions from the teachers were in the form of rules.

Introduction

Consider listening to your very first speech by a politician. Your previous knowledge is likely to influence your attitude towards her or him. Perhaps, already from childhood your father has imprinted in you that politicians are guided by strictly egoist motives and your general conceptions thus include a belief that no politician advocates a proposal that does not lay in his or her personal interest. You hear a short speech that is neutral in content. Later you listen to another politician. How is your opinion of this second politician influenced by the first encounter? You did not receive much useful feedback from the first encounter, as you only listened to a short neutral speech. However, in your memory the first politician is stored as a person only interested in pursuing his own interest. This exemplar memory only in part derives from direct experience with a politician; in part it is derivative of more general beliefs held prior to the encounter. However, by now this belief is supported also by "concrete experience" of politicians.

This was an example of a real world situation that is different from most categorization experiments where classification models are tested in tasks with simple perceptual stimuli and trial-by-trial outcome feedback. Everyday situations often do not offer direct and unambiguous feedback and exemplar memory is thus likely to derive in part also from other sources of knowledge besides concrete experience with the objects.

There has been increasing interest in multiple representation levels (e.g., Ashby, Alfonso-Reese, Turken, & Waldron, 1998) and there is evidence that people can adaptively shift between different representation levels in response to the experimental demands (Jones, Juslin, Olsson, & Winman, 2000). With experience, knowledge first represented abstractly may be projected onto concrete exemplars so that in the end the beliefs are supported also by an extensive exemplar memory: a phenomenon that might be called *pseudo-experience*.

Even if some extensions of exemplar models allows for storage of exemplars as they are interpreted and not solely in terms of their physical properties (e.g., the model presented by Smith & Zarate, 1992), the argument supporting this claim is based on general observations and not linked to predictions from different models. For example, one such observation is that a reencounter of a stimulus facilitate the same reactions and processes (see the review in Smith & Zarate, 1992).

In this paper, we examine the possibility of extending the scope of *exemplar models* (Medin & Schaffer, 1978; Nosofsky, 1986) to situations where people do not receive outcome feedback, but form beliefs about the criterion from abstract knowledge of rules. In these circumstances, one possibility is that people completely abandon exemplar processes as a basis for their judgments and directly use abstract knowledge in the form of rules or prototypes.

Another possibility is that people generate the criteria from abstract knowledge and store them together with the experienced exemplars; later to rely on these stored exemplars to make their judgments. We explore these possibilities in a "teacher-student" task where a teacher learns how to judge the toxicity of bugs from outcome *feedback* and the student has to rely on *feedforward* summary information provided by the teacher. The question is if there is evidence for exemplar processing in the students judgments even if they do not receive feedback or instructions about exemplars from the teachers.

Measuring Exemplar Effects

To develop an exemplar effect index we need to consider a category structure that allows us to differentiate between predictions by the exemplar model and other plausible models, in this case a cue-abstraction model that linearly integrates cues. The results previously obtained with this task revealed large individual differences and a shift from exemplar memory to more mental cue-integration processes when the criterion is changed from classification to a continuous judgment task (Juslin, Olsson, & Olsson, 2002).

The task requires participants to use four binary cues to infer a continuous criterion. (Juslin et al., 2002). The judgments involve the toxicity of subspecies of a fictitious Death Bug. The different subspecies vary in concentration of poison from 50 ppm (harmless) to 60 ppm (lethal). The toxicity can be inferred from four visual cues of the subspecies (e.g., the length of their legs, color of their back).

The binary cues C_1, C_2, C_3, and C_4 take on values 1 or 0. The toxicity c of a subspecies is a linear, additive function of the cue values:

$$c = 50 + 4 \cdot C_1 + 3 \cdot C_2 + 2 \cdot C_3 + 1 \cdot C_4 \cdot \tag{1}$$

C_1 is the most important cue with *coefficient* 4 (i.e., a relative weight .4), C_2 is the second to most important with coefficient 3, and so forth. A subspecies with feature vector (0, 0, 0, 0) thus has poison concentration 50 ppm; a subspecies with feature vector (1, 1, 1, 1) has 60 ppm. The continuous criteria for all 16 subspecies (i.e., possible cue configurations) are summarized in Table 1.

Table 1
Structure of the Task

Exemplar	C_1	C_2	C_3	C_4	Criterion	Exemplartype
1	1	1	1	1	60	E
2	1	1	1	0	59	T
3	1	1	0	1	58	T
4	1	1	0	0	57	O
5	1	0	1	1	57	N
6	1	0	1	0	56	N
7	1	0	0	1	55	N
8	1	0	0	0	54	T
9	0	1	1	1	56	O
10	0	1	1	0	55	O
11	0	1	0	1	54	T
12	0	1	0	0	53	T
13	0	0	1	1	53	T
14	0	0	1	0	52	T
15	0	0	0	1	51	T
16	0	0	0	0	50	E

Note: C = Cue; E = Extrapolation; T = training exemplar; O = Old comparison exemplar in training, N = New comparison exemplar presented at test.

In *training*, the participants encounter 11 subspecies and make *continuous judgments* about the toxicity of each subspecies ("The amount of poison is 57%"). As indicated in the two right-most columns of Table 1, five subspecies are omitted in training. In a *test phase*, the participants make the same judgments as in the training phase, but for all the 16 subspecies and without feedback. The task allows perfect performance in training both by exemplar memory and induction of the task structure (i.e., by inducing the cue weights in Eq. 1).

The *exemplar model* implies that participants make judgments by retrieving similar exemplars (subspecies) from long-term memory. The *context model* of classification (Medin & Schaffer, 1978) suggests that in a task that only requires participants to judge if a bug is dangerous or not, the probability $p_E(\hat{b}=1)$ of categorization as dangerous (1) equals the ratio between the summed similarity of the judgment probe to the dangerous exemplars and the summed similarity to all exemplars:

$$p_E(\hat{b}=1) = \frac{\sum_{j=1}^{J} S(p, x_j) \cdot b_j}{\sum_{j=1}^{J} S(p, x_j)}, \tag{4}$$

where p is the probe to be judged, x_j is stored exemplar j (j=1…J), $S(p,x_j)$ is the similarity between the probe p and exemplar x_j, and b_j is the binary criterion stored with exemplar j (b_j=1 for dangerous, b_j=1 for harmless). J depends on the size of training set of exemplars.

The similarity between probe p and exemplar x_j is computed by the multiplicative similarity rule of the context model (Medin & Schaffer, 1978):

$$S(p, x_j) = \prod_{i=1}^{4} d_i, \tag{5}$$

where d_i is an index that takes value 1 if the cue values on cue dimension i coincide (i.e., both are 0 or both are 1), and s_i if they deviate (i.e., one is 0, the other is 1). s_i are four parameters in the interval [0, 1] that capture the impact of deviating cues (features) on the overall perceived similarity $S(p,x_j)$. s_i close to1 implies that a deviating feature on this cue dimension has no impact at all on the perceived similarity and is considered irrelevant. s_i close to 0 means that the overall similarity $S(p,x_j)$ is close to 0 if this feature is deviating, assigning crucial importance to the feature. The parameters s_i capture the similarity relations between stimuli and the attention paid to each cue dimension, where a lower s_i signifies higher attention.

The context model was developed for classification, in most cases to binary categories. To generate predictions also for judgments of a continuous criterion we relax the model by allowing the outcome index b_j to take not only binary but also continuous values. The estimate $\hat{c}_E$ of the criterion c is a weighted average of the criteria c_j stored for the exemplars, where the similarities $S(p,x_j)$ are weights (see e.g., Juslin & Persson, 2000; Smith & Zarate, 1992, for similar applications).

The *cue-abstraction model* assumes that the participants abstract explicit cue-criterion relations in training, which are mentally integrated at the time of judgment. When presented with a probe the participants retrieve rules connecting cues to the criterion from memory (e.g., "Green back goes with being poisonous"). The rules specify the sign of the relation and the importance of each cue with a cue weight. For example, after training the rule for cue C_1 may specify that C_1=1 goes with a large increase in the toxicity of a subspecies.

Cue abstraction implies that participants compute an estimate $\hat{c}_R$ of the continuous criterion c. For each cue, the appropriate rule is retrieved and the estimate of c is adjusted according to the cue weight ω_i (i=1…4). The final estimate $\hat{c}_R$ of c is a linear additive function of the cue values C_i,

$$\hat{c}_R = k + \sum_{i=1}^{4} \omega_i \cdot C_i \,, \qquad (2)$$

where $k = 50 + .5 \cdot (10 - \sum \omega_i)$. If ω_1=4, ω_2=3, ω_3=2, and ω_4=1, Eq's 1 and 2 are identical and the model produce perfect judgments. The intercept k constrains the function relating judgments to criteria to be regressive around the midpoint (55) of the interval [50, 60] specified by the task instructions.

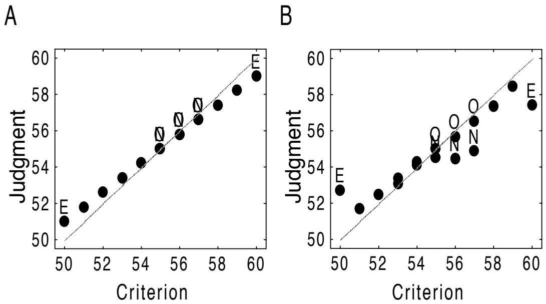

Figure 1*:* Panel A: Cue abstraction model predictions for the constrained training set. Panel B: Exemplar model predictions with similarity parameter s=.1 for the constrained training set. O = Old comparison exemplar in training, N = New comparison exemplar presented at test, E = Extrapolation exemplars.

The predictions from the two models are summarized in Figure 1 and shows that the models predict distinct patterns. The cue abstraction model allows accurate extrapolation beyond the distribution of criteria in the training set [51, 59]. Whenever the correct signs of the cue weights are identified, the most extreme judgments are made for exemplars 1 and 16. The exemplar model, that computes a weighted average of the criteria observed in training, can never produces a judgment outside the observed range. The most extreme judgments are made for criteria c=51 and 59.

Moreover, with cue abstraction there should be no systematic differences between judgments for “New” and “Old” exemplars with c=55, 56, and 57: the process is the same in both cases. However, with the exemplar model there is more accurate judgments for Old exemplars, because these judgments benefit from retrieval of identical exemplars with the correct criterion.

These differences in predictions allow us to define measures of the amount of exemplar processing. First, The *old-new difference index* ΔON is defined as,

$$\Delta ON = \bar{d}_{Old} - \bar{d}_{New} \qquad (6)$$

where $\bar{d}_{Old}$ is the mean absolute deviation between judgment and criterion for the three old exemplars and $\bar{d}_{New}$ is the corresponding mean deviation for the three new exemplars in Table 1. When judgments for old rather than new exemplars are more accurate, the index is negative. The *extrapolation index EI* is the mean deviation from linear extrapolation,

$$EI = \frac{(x_{51} - x_{50}) + (x_{60} - x_{59}) - 2b)}{2} \qquad (7)$$

where x_{50}, x_{51}, x_{59}, and x_{60} are the judgments for exemplars with criteria 50, 51, 59 and 60, respectively. The value of b is determined by the difference x_{51} - x_{50} (or equivalently x_{60} - x_{59}) predicted by a linear regression relating judgments to criteria. Perfect linear extrapolation implies an extrapolation index that is 0 (e.g., when the judgments are perfectly accurate). If the index is negative, the exemplars with extreme criteria do not receive as extreme judgments as implied by linear extrapolation. For example, the indices in Figures 1A are 0, but the indices in Figure 1B are negative. The mean of ΔON and EI provides a total index of exemplar effects, *Total EE*.

Method

Participants

Forty undergraduate students participated. The participants were paid between 50 and 100 SEK depending on their performance.

Apparatus and Materials

The experiment was carried out on a PC–compatible computer. The exemplars varied in terms of four binary cues; leg length (short or long), nose length (short or long), spots or no spots on the fore back and two patterns on the back. The cues and the cue values in the abstract structure in Table 1 were randomly assigned to new concrete visual features for each new participant. Two types of presentation modes were used, one *analogue* where drawings of the bugs were presented and one *propositional* with written descriptions of the bugs.

Design and Procedure

The task was done in pairs, one participant in each pair was randomly assigned as teacher and the other as student. Half of the teacher-student pairs were randomly assigned to the analogue condition and the other half to the propositional condition.

The written instructions informed the participants that the task involved judgments of the toxicity of subspecies of a Death bug from 50 to 60 ppm and that the difference between teachers and students was that the teacher receives outcome feedback and that the student does not. The participants were also informed that they would receive a minimum payment of 50 SEK and up to 100 SEK *depending on the performance of the student*. The performance bonus was calculated by taking half the correlation between the students’ judgments and the criterion values in the test phase times 100. In addition, the teacher was told that after each training block they were to write down instructions to the student on how to assess the toxicity of the bugs as a Word file. The teachers were free to give any instructions they wanted. The word files were collected for further analysis by the experimenter. No additional contacts between teachers and students were allowed (except for strictly clarifying questions in regard to spelling errors, as mediated by the experimenter).

The training phase consisted of four blocks with 55 trials each making a total of 220 trials. Only the teacher received outcome feedback of the correct toxicity level in the learning phase. After each training block the teacher wrote down instructions to the student and the experimenter handed it over to the student that were seated in another room. In the test phase each exemplar were presented four times without feedback for both teachers and students making a total of 64 trials. The entire experiment took from one hour and fifteen minutes to two hours.

Results

Performance for teachers and students in the learning phase of the analogue and the propositional conditions measured by the absolute deviation between judgment and criterion are presented in Figure 2. The performance for teachers and students are about the same in the last part of training. Performance is better in the analogue condition than in the propositional condition.

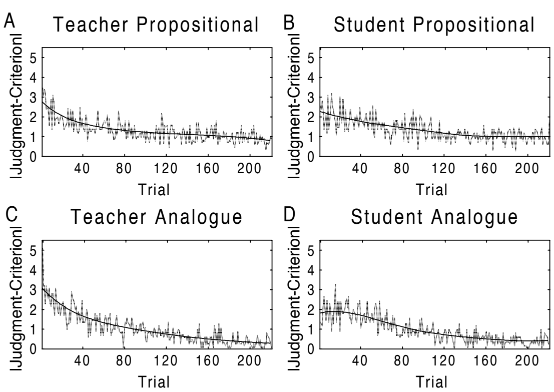

Figure 2: Panels A-D: Mean absolute deviation between judgment and criterion for teachers and students. The curves are fitted according to a negative exponentially weighted smoothing procedure.

Figure 3 shows that there appears to be some exemplar effects in all the conditions, for both teachers and students. The judgments of the new exemplars are less accurate than the old exemplars and underestimate the criteria. The figure also shows that both teachers and students have some difficulties with extrapolating beyond the distribution of criteria in the training set.

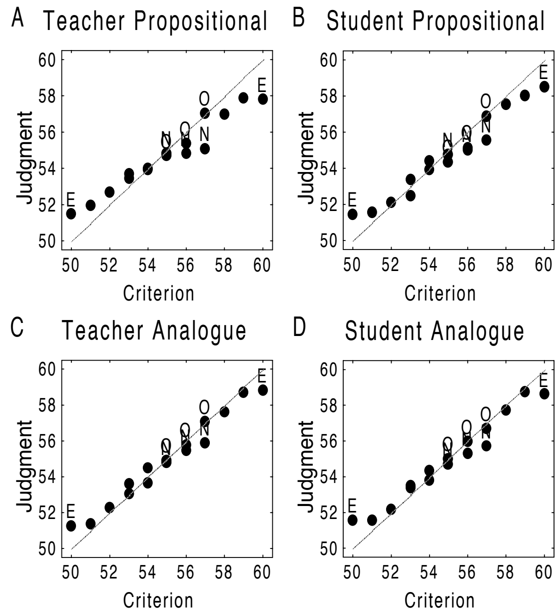

Figure 3: Panels A-D: Mean judgments for the different criteria for teachers and students.

The exemplar indexes were collapsed over the analogue and the propositional conditions as no significant differences were found between the two conditions for any of the indexes. Figure 4 shows that there are clear exemplar effects for both teachers and students, as the 95% confidence intervals does not include zero.

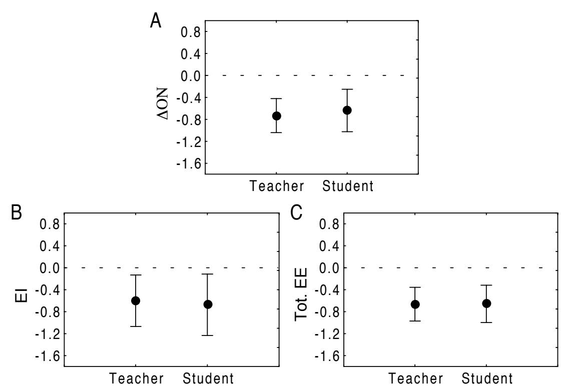

Figure 4: Panel A: Mean Old-New difference index, ΔON. Panel B: Mean *extrapolation index EI*. Panel C: Mean total exemplar effects index, *EE*. The error bars are 95% confidence intervals.

We coded the instructions the teachers gave the students as containing exemplars or not by a strict coding scheme that assigned any ambiguous case as an exemplar instruction. Six teachers had instructions containing exemplar information, for example "Green body, short legs, long nose, no spots = 51%". A typical part of an instruction that did not contain exemplar instruction was "Begin with 50% and add: short grey legs +2…green long legs +0 [and so on for all the cues]".

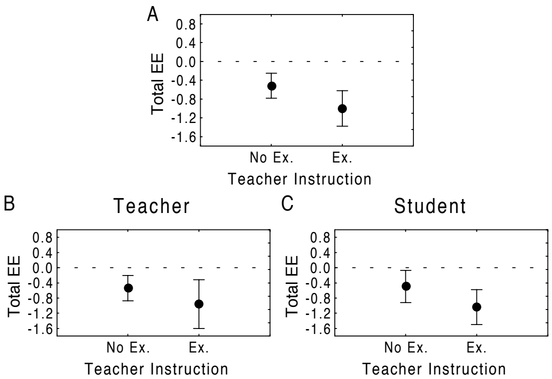

Figure 5: Panel A: Mean total exemplar effects index, *EE*, for No exemplar instructions and Exemplar instructions over all conditions. Panel B: Mean total exemplar effects index, *EE*, for No exemplar instructions and Exemplar instructions for teachers. Panel C: Mean total exemplar effects index, *EE*, for No exemplar instructions and Exemplar instructions for students. The error bars are 95% confidence intervals.

Shown in Figure 5 is the total exemplar effects index separately for the participants with no exemplars in the teachers' instructions and those with exemplars in the instruction. It can be seen that there is an effect of instruction with larger exemplar effects for exemplar instructions, $t(38) = 2.00$, $p = .026$, one-tailed. More importantly, students with no exemplar instructions from their teachers exhibit evidence for exemplar processing, as the confidence intervals do not include zero.

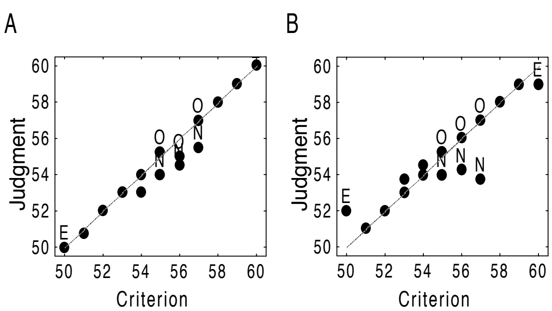

Figure 6: Mean judgments for two students that did not receive any exemplar information. Panel A: A student better described with a cue-abstraction model. Panel B: A student better described with an exemplar model.

Investigation of individual participants data reveals large individual differences. Figure 6 shows two students that did not receive any exemplar information from their teachers. One student is better described as a cue abstraction participant and the other student as an exemplar participant.

Discussion

In this paper, we have shown that people can spontaneously rely on pseudo-experience in a judgment task. Even when the information people receive does not contain information about specific exemplars, people cannot help project abstract rule knowledge onto concrete exemplars and then use these exemplars in the judgment process.

Even if there are large individual differences in data with some people operating only in accordance with the cue-abstraction model, it seems difficult for most people to totally abandon exemplar processing. Even if you initially execute a rule to determine what response you will make, the very act of executing the rule implies processing the exemplar in front of you. For example, you need to scan the object for features that fit the rule conditions. Even if you do not consciously trying to remember exemplars, the end result is incidental learning of exemplars that later influences judgments.

One caveat is that the types of categories used could affect the prevalence of exemplar based processing. For example, the results in a series of experiments by Minda, Smith and colleagues (e.g., Minda & Smith, 2001) suggest that larger categories, better structured categories, and more complex stimulus promotes prototype processing at the expense of exemplar processing.

In this, and other tasks, rule based representations provide great powers of generalization and communication. One answer why we sometimes cannot avoid storing and using exemplars may be found in the idea that our cognitive system consists of multiple levels of representation that work together or compete to determine responses in specific tasks (see e.g., Ashby et al., 1998). Our results fit into the notion that the exemplar representation has a function that separates it from other representational formats in that it acts like an automatically activated back-up system that preserves distributional and individuating information about the world.

Acknowledgments

Bank of Sweden Tercentenary Foundation supported this research.

References

Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review, 105,* 442-481.

Jones, S., Juslin, P., Olsson, H., & Winman, A. (2000). Algorithm, heuristic or exemplar*:* Processes and representation in multiple-cue judgment. In L. Gleitman, & A. K. Joshi (Eds.), *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society* (pp. 244-249). Hillsdale, NJ: Erlbaum.

Juslin, P., Olsson, H., & Olsson, A-C. (2002). *Abstract and concrete knowledge in categorization and multiple-cue judgment*. Manuscript submitted for publication. Department of Psychology, Umeå University, Umeå, Sweden.

Juslin, P., & Persson, M. (2000). *PROBabilities from EXemplars (PROBEX): A "lazy" algorithm for probabilistic inference from generic knowledge*. Manuscript submitted for publication. Department of Psychology, Umeå University, Umeå, Sweden.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review, 85,* 207-238.

Minda, J. P., & Smith, J. D. (2001). Prototypes in category learning: The effects of category size, category structure, and stimulus complexity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27,* 775-799.

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

Smith, E. R., & Zarate, M. A. (1992). Exemplar-based model of social judgment. *Psychological Review, 99,* 3-21.