Cogn Process
DOI 10.1007/s10339-008-0243-x

RESEARCH REPORT

Category learning from equivalence constraints

Rubi Hammer · Tomer Hertz · Shaul Hochstein · Daphna Weinshall

Received: 31 January 2007 / Revised: 5 July 2007 / Accepted: 3 November 2008

Abstract Information for category learning may be provided as positive or negative equivalence constraints (PEC/NEC)—indicating that some exemplars belong to the same or different categories. To investigate categorization strategies, we studied category learning from each type of constraint separately, using a simple rule-based task. We found that participants use PECs differently than NECs, even when these provide the same amount of information. With informative PECs, categorization was rapid, reasonably accurate and uniform across participants. With informative NECs, performance was rapid and highly accurate for only some participants. When given directions, all participants reached high-performance levels with NECs, but the use of PECs remained unchanged. These results suggest that people may use PECs intuitively, but not perfectly. In contrast, using informative NECs enables a potentially more accurate categorization strategy, but a less natural, one which many participants initially fail to implement—even in this simplified setting.

R. Hammer (✉) · T. Hertz · S. Hochstein · D. Weinshall
Interdisciplinary Center for Neural Computation,
Hebrew University, Edmond Safra Campus,
91904 Jerusalem, Israel
e-mail: rubih@alice.nc.huji.ac.il

R. Hammer · S. Hochstein
Neurobiology Department, Institute of Life Sciences,
Hebrew University, Jerusalem, Israel

T. Hertz · D. Weinshall
School of Computer Sciences and Engineering,
Hebrew University, Jerusalem, Israel

Keywords Category learning · Categorization · Concept acquisition · Dimension weighting · Perceived similarity · Rule-based · Learning to learn

Introduction

Since the early days of cognitive research, a number of theories have been suggested to describe both the structure of categories and the mental processes involved in their acquisition. The *classical view* suggests that categories may be described by a list of necessary and sufficient attributes that determine category membership (e.g. Katz and Postal 1964; Smith and Medin 1981). Similar ideas are still prevalent for category learning tasks where categories can be described with an explicit rule (Shepard et al. 1961; Mooney 1993; see also Ashby and Maddox 2005 for recent view). On the other hand, *probabilistic theories* suggest that objects are categorized by similarity to an internal representation of a category prototype (Rosch and Mervis 1975) or category exemplars (Medin and Schaffer 1978; Nosofsky 1987, 1988, 1990). As an object's similarity to this representation increases, the probability that it belongs to the represented category also increases.

A common theme in most of these views is that the process of category learning requires learning about the relevance of specific object properties for categorization (Rouder and Ratcliff 2006). Some views of the role of similarity in categorization explicitly take this issue into consideration, suggesting that objects are grouped together based on their similarity in specific features (Tversky 1977; Tversky and Gati 1982) or within specific feature-dimensions perceived as more relevant for categorization (Garner 1978; Nosofsky 1987; Medin et al. 1993; Goldstone 1994a).

Published online: 03 December 2008

Apparently, evaluating the importance of different object properties is essential for category learning. In the current study we take a novel approach to address this issue of "dimension weighting" in category learning. We show that dimension weighting can be learned from a training set of equivalence constraints, which indicate the pair-wise relationship between exemplars. We provide both conceptual and empirical evidence for an asymmetry in the contributions of two types of equivalence constraints and demonstrate the implications of this asymmetry for numerous categorization scenarios.

Category learning from equivalence constraints

We call a restriction indicating that two exemplars belong to the same category a positive equivalence constraint (PEC), and a restriction indicating that two exemplars belong to different categories a negative equivalence constraint (NEC). We claim that rule learning or dimension weighting can be performed naturally when a classifier is provided with equivalence constraints. Since equivalence constraints can be used for extracting a rule or to restrict the perception and/or use of similarities between objects within a category, or dissimilarities between objects of different categories, classifiers can generalize from constrained examples to other objects encountered later.

Both PECs and NECs are available in a variety of category learning scenarios. For example, when a parent tells a child—pointing to animals unfamiliar to the child—"This is a dog and that is also a dog," he or she indicates to the child that the two animals belong to the same category. When the parent then labels two other animals as "These are horses," he or she provides the child not only with an indication that these two belong to a single category, but also that the latter two animals differ from dogs and belong to a different category. Here, labels are used for identifying relations between exemplars, and, as is often the case with labels, the information provided mixes PECs and NECs.

Naïve participants performing a supervised same/different task, where labels are not provided, also learn relationships among a few objects. When the participant guesses that two objects belong to the same category (or to different categories), feedback provided by a supervisor, indicating that he was right or wrong, ultimately provides him with an indication of whether the two truly belong to the same category or to two different categories (e.g. Cohen and Nosofsky 2000; Goldstone 1994b). Similar principles underlie supervised categorization tasks in which stimuli from two or more different categories are presented in sequential order see (Ohl et al. 2001, for an example, of a category learning task with animals participating in a go/no-go paradigm). Similarly, in everyday scenarios, when a child asks an adult, "Is this one the same as that one?" a yes/no response indicates whether the two are from the same category or from different categories.

This principle of category learning from indications that some objects are or are not from the same category is not limited to category learning with explicit supervision. Many contextual clues can indicate whether objects are from the same category or from different categories. For example, seeing two animals playing together, one may assume that they are from the same species, while seeing one animal chasing another may indicate that the two are not the same. Such scenarios provide clues to the relations among objects without direct supervision and may contribute to category learning as much as scenarios in which direct supervision is available. In fact, current approaches to category learning argue that categories can be learned and representations built from acquired relations among exemplars (Gentner and Kurtz 2005; Jones and Love 2004).

In the current study we tested category learning when providing participants with either only PECs or only NECs. We used a rule-based categorization task in which stimuli were defined by five binary dimensions (see Allen and Brooks 1991 for a similar stimulus design). To maximize performance, participants had to identify two or three relevant dimensions in each categorization task (experimental trial). Guided by the studies reviewed above, we believe that using a rule-based task is a plausible approach for studying the role of equivalence constraints in dimension weighting for category learning. Furthermore, using this simple setup enables us to control the amount of information provided by each of the two types of constraints.

The motivation for evaluating the separate contributions of PECs and NECs for category learning arises from the differences between these two types of constraints: NECs are more common than PECs, but generally PECs are more informative than NECs; PECs specify within-category variation, while NECs specify between-category variation; and PECs are transitive, but NECs are not. In the next section we survey these differences in detail, pointing out the importance of these differences for dimension weighting in the context of category learning.

Differences between positive and negative equivalence constraints

NECs are more common than PECs

In most natural scenarios NECs abound while PECs are less common. This simple observation is demonstrated in Fig. 1. Figure 1a presents a natural scene with only three animal categories (5 antelopes, 3 giraffes and 3 zebras). The number of PECs is the number of possible pairs of

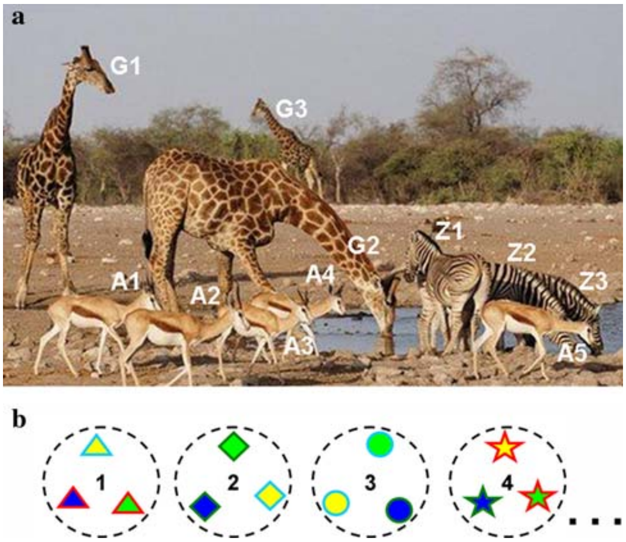

Fig. 1 Demonstration of difference between PECs and NECs in terms of availability. **a** A natural scene with three different categories, including 16 PECs and 39 NECs (see text). **b** An illustration in which each *dashed circle* represents a category of three objects (*inner shapes*); there are four categories, 12 PECs, and 54 NECs (see text). Generally, in any scenario representing three or more objects taken from two or more categories, the number of NECs will always be higher than the number of PECs (see text and Appendix)

antelopes (10), giraffes (3) and zebras (3), for a total of 16 PECs in all. The number of NECs is the number of possible pairs composed of two animals from two different categories, (15 antelope-giraffe, 15 antelope-zebra and 9 giraffe-zebra pairs), which is 39 NECs in total. Thus, the difference between the number of NECs and PECs is large even in a scene with only three categories.

If we add more categories, the number of NECs increases more dramatically than the number of PECs, as illustrated in Fig. 1b. Here, each category (1–4) is composed of only three objects. The number of PECs in each category is three. In a world which includes only categories 1 and 2, there are six PECs and nine NECs. When we include category 3, the number of PECs increases by 3 and the number of NECs by 18. When category 4 is added, the number of PECs increases again by 3, reaching 12, but the number of NECs is doubled from 27 to 54. The general rule is that when more categories are added, the number of PECs increases linearly while the number of NECs increases as a quadratic polynomial; (see Appendix 1 for a formal proof).

Between-Category versus within-category variations

Both PECs and NECs may play an important role in identifying features or dimensions which enable grouping objects into categories or discriminating between categories. Yet, the two types of constraints differ: When we learn that two novel objects are from the same category (PEC), we can expect that at least some of the dimensions for which the two objects share similar values (features) are relevant to categorization. More definitively, we can conclude that *all* dimensions discriminating between the two objects are generally irrelevant, and that these differences only reflect within-category variation along these dimensions. Thus, the amount of information provided by a PEC is related to the number of irrelevant dimensions that it indicates.

The case of NECs is more complex: When we are told that two objects are from different categories, if the objects differ by more than one dimension—which is the case for most NECs—then we cannot definitively conclude which of these dimensions is relevant for discriminating between the categories. In fact, a salient non-relevant dimension may mask detection of a relevant less salient dimension (e.g. Huettel and Lockhead 1999). Similarly, we cannot determine whether a dimension in which the two objects share the same value is relevant or irrelevant for categorization, since two objects from different categories may share many features, as long as they differ by at least one feature that is relevant for categorization. The only time we can confidently learn which dimension is relevant when provided with a NEC, is when there is only a single dimension by which the negatively constrained objects can be discriminated (see Goldstone 1994b for related ideas). In this case, we can conclude that this unique discriminating dimension is necessarily relevant for categorization. At the same time, even in this special case, we cannot confidently infer anything about the relevance or irrelevance of the other dimensions.

PECs and NECs do not provide the same amount of information

As we have seen earlier (subsection 1), NECs are much more common than PECs and therefore they might be expected to be a more readily available source of information in most scenarios. However, as already implied (subsection 2), the reverse is true: despite their greater number, most NECs are only poorly informative for the task of identifying relevant dimensions. Since the amount of information provided by a PEC depends on the number of irrelevant dimensions it specifies, while a NEC at best specifies one dimension as relevant, we conclude that most PECs provide more than one bit of information while NECs provide at the most one bit of information, and that too rarely.

An example is shown in Fig. 2. Assume that objects A, B, C, D belong to category 1, while E, F, G, H belong to category 2. Basically, the two categories can be discriminated by object color but not by texture or shape, where

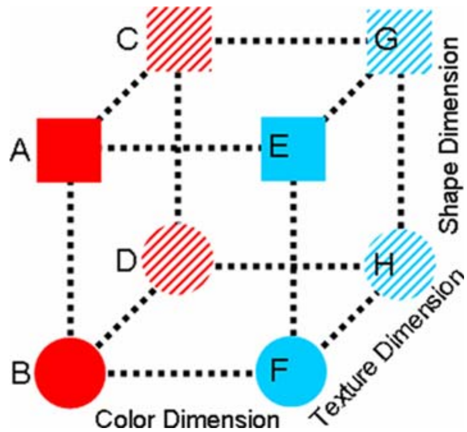

Fig. 2 Example of a three-dimensional object space with two categories. In this simplified example each dimension is binary (i.e., has only two values/features). The dimensions are color (*red* vs. *blue*) shape (*circle* vs. *square*) and texture (*filled* vs. *dashed*)

differences are due to within-category variation. Our naïve classifier is not aware of this category setting or the number of existing categories that will have to be learned from the given constraints.[1]

In this setting there are 12 possible PECs that can be given to the classifier: {A = B, A = C, A = D, B = C, B = D, C = D} from category 1, and {E = F, E = G, E = H, F = G, F = H, G = H} from category 2. All PECs are informative to some extent. For example, from being informed that A = C or F = H, the classifier can learn that texture is not a relevant dimension. Such PECs provide one bit of information (decisively indicating that one dimension is not relevant). From A = D or E = H, one can learn that both texture and shape are not relevant. Such PECs provide two bits of information. When provided with one of the latter PECs only the color dimension is left with a potential discriminating value enabling discrimination between the assumed categories. The probability that the classifier can extract this information from only two different random-lyselected PECs is high. For example, in Fig. 2, where each category includes four different objects, the probability that two randomly selected positive constraints identify the two irrelevant dimensions is $P = 1 - 2 \times 4/12 \times 3/11 \approx 0.82$.

The number of NECs (16) is greater than the number of PECs (12), but only four are useful in definitively identifying a relevant dimension. For instance, if told that A ≠ G, the classifier cannot learn whether the two objects belong to two different categories because they differ in color, in texture or in both. He or she also cannot conclude whether shape is relevant for discriminating between categories, since shape could be one of two determinants. The only way to determine that color is the critical discriminating dimension between categories is to be provided with one of the NECs which present two objects that differ solely in color: A ≠ E, C ≠ G, B ≠ F or D ≠ H. Each one of these NECs provides one bit of information. With only one relevant dimension, the probability that at least one of two randomly selected NECs will point to the relevant dimension is $P = 1 - 12/16 \times 11/15 = 0.45$ Note, in addition, that participants can only derive that a dimension is irrelevant from the absence of NECs demonstrating that it is relevant.

In the aforementioned example, only one of the existing three dimensions was relevant. When the number of relevant dimensions is increased, the number of categories is also increased; (e.g. for binary dimensions and a conjunctive classification rule, each added relevant dimension doubles the number of categories). With this increase, the chance of learning which dimensions are relevant from randomly selected NECs is dramatically reduced. At the same time PECs remain highly informative.

For example, if both color and shape are relevant for categorization, our object space is now divided into four categories: {A, C}, {E, G}, {B, D}, and {F, H}. The number of PECs is now reduced to four (one in each category), but all of them indicate that texture is irrelevant for categorization (due to the within-category variation in texture). We can therefore identify the irrelevant dimension from each one of the four PECs, and (assuming that all dimensions that have not been shown to be irrelevant are indeed relevant) the probability of learning the remaining relevant dimensions from two PECs is 1. In contrast, while the number of NECs has increased to 24, only 8 of them are informative by negatively constraining pairs that differ in only one of the two relevant dimensions. The probability of learning the two relevant dimensions from two *essentially* different NECs is now quite low: $P = 2 \times 4/24 \times 4/23 \approx 0.058$ (roughly 6%). It seems that despite NECs being much more common than PECs, retrieving valuable information from them may be as wearisome as separating the wheat from the chaff.

The role of transitivity

Another important property differentiating PECs and NECs is transitivity. Whereas PECs are transitive, NECs are not. This property of PECs is expected to be quite helpful in the context of category learning. Using transitivity, PECs can help in "packing together" objects into categories. For example, by being informed that A and B are from the same category, and that C and D are from the same category, it is enough to be also informed that A and C are from

[1] For clarity of presentation and simplicity of experimentation, this example—as our experimental paradigm—uses binary feature values and categories defined by rules. Nevertheless, the conclusions of the analysis—as the results of the study—extend to other categorization scenarios.

the same category to be able to “pack” the four objects together into one category. As the number of objects increases, the contribution of transitivity in “packing objects” into categories also increases. This property of PECs can be helpful in summing together pieces of information when constructing an internal representation of newly learned categories. For example, transitivity may facilitate forming a category prototype by averaging the common relevant features of the packed-together objects. Similarly, a small subset of packed-together objects can be used as a set of exemplars representing the category.

Since NECs are not transitive, it is much harder to accumulate information for a categorization task when using only NECs. In the current study we will not directly address transitivity. We will only assume that this property of PECs may bias people in their use of the two types of constraints even in a task in which transitivity is neutralized.

Summing up the differences between PECs and NECs This theoretical overview has highlighted a number of inherent differences between PECs and NECs, and pointed toward a weakness of NECs: Whereas all PECs are informative for indicating irrelevant dimensions, NECs are informative in decisively indicating a relevant dimension only on those rare occasions when we are informed that two objects that are similar in most of their properties, nevertheless belong to two different categories. The second disadvantage of NECs when compared to PECs, (which is less important for the current experimental design), is that PECs are transitive whereas NECs are not.

One could claim that a supervisor, when accessible, could save classifier effort by providing more useful PECs and NECs. For example, Avrahami et al. 1997 showed that in some cases “expert participants” teach “novices” new categories by using sequences of exemplars that identify the borders on each specific dimension, demonstrating members of the target category and non-category examples. However, in natural scenes this selection might be difficult, since, as we have seen, informative NECs are rare; in addition, direct and explicit guidance may not always be available.

Earlier findings also implied that there are differences in the way people use PECs and NECs. It was found that when participants are asked to define a target category using selected exemplars, they were biased toward using positive examples but not negative examples—i.e., defining the target category using only member stimuli, avoiding use of stimuli from outside the category (Wason 1960; Klayman and Ha 1987). This bias enables comparing stimuli within a category (PECs) but not comparing stimuli between categories (NECs). These findings are puzzling when faced with evidence demonstrating an advantage of using both positive and negative examples in categorization tasks (Levine 1966) or when learning a mathematical rule (Kareev and Avrahami 1995).

Further research revealed that when a target category could be defined by a simple rule, people favored using positive examples, concentrating on “positive-ideal” stimuli from the target category that were relatively far from a category border. This selection is best for identification of task-irrelevant dimensions. On the other hand, when the border defining the target category was a diagonal line integrating two dimensions, more participants used both positive and negative examples, including stimuli that were close to the border from both sides (Avrahami et al. 1997). Similarly, in cases of exemplar-based categorization, a more refined representation is achieved when using both similarities between same-category exemplars and dissimilarities between different-category exemplars (Stewart and Brown 2005). These two lines of findings are consistent with the idea that more confusable categories—such as medical diagnosis of similar syndromes—require a representation that enables their comparison (Brooks et al. 1991; Kulatunga-Moruzi et al. 2001).

A difference in the use of PECs and NECs was also found in a category learning task with sequential-presentation of training stimuli. When participants were provided with a sequence of items from the same category, classification of a novel item was easier than if the provided examples were from different categories (Whitman and Garner 1962). Recent sequential-presentation studies also imply that the categorical relation between presented exemplars in a sequence affect the way later items are categorized (Jones et al. 2006; Stewart et al. 2005). Nevertheless the goal of these experiments was to study the effect of perceptual and cognitive factors, such as memory and contrast, and they do not provide a direct evaluation of differential PEC vs. NEC contributions.

In summary, since PECs and NECs provide the classifier with potentially different types of insight and since there is evidence implying that they are used differently, it becomes important to directly investigate how people use these two types of constraints and to what extent the two types are integrated in categorization tasks. While the current research uses binary and discrete feature-dimensions, the implications regarding information provided by PECs versus NECs are relevant whenever dimension weighting is involved. Similarly, in our experimental setting constraints are definitive regarding dimension relevance or irrelevance (binary weight of 0 or 1), but more refined dimension weights could be achieved by using a larger number of constraints, (e.g. with many indicating one dimension as relevant, implying a high weight, and fewer indicating another as relevant, implying low weight). We believe that analyzing the separate contributions of these two building blocks of category learning can provide useful insights for understanding categorization errors, and may shed light on a number of known phenomena in category learning, such as the related findings described above.

Outline of the experiments and their motivation

In the current study, in each one of the ten experimental trials, participants performed a different categorization task in which they used exclusively either PECs or NECs for identifying the task relevant dimensions. The first experiment tested performance with randomly selected PECs or NECs. Results confirmed our prediction that performance is better with PECs than with NECs. However, recall that this prediction derived from the fact that typical PECs provide more information than typical NECs. Thus, this result could reflect simply the information provided by the constraints and not the proficiency of their use by the participants.

Experiment 2 therefore tested the use of PECs and NECs when these are specifically chosen to provide the same amount of information. Importantly, we find a difference here, too, in the performance with PECs versus NECs. This difference must reflect the use of these constraints, rather than their inherent information content. Interestingly, we find that people may be divided into two groups—those who are able to use NECs quite well, and those who are unable to do so. This raises the possibility that using NECs is non-intuitive and that it is difficult for some to derive the proper strategy for their use. Therefore, in Experiment 3, we provided all participants with directions for the use of either PECs or NECs. Here we find that all participants succeed in the use of either type of constraint, supporting the prediction that the difference between PECs and NECs in natural circumstances leads to different proficiencies in their use.

Experiment 1: baseline performance

The first experiment was designed to measure baseline performance. As in all our experiments, categories were defined by the conjunction of their features along two or three relevant dimensions. In this experiment there were three experimental conditions: in the first, participants categorized stimuli when no Equivalence Constraints were provided to them (the noEC condition). This condition was needed to assess the contribution of equivalence constraints that were provided in the other experimental conditions. In the second and third experimental conditions, participants were provided with randomly generated positive (randPEC) or negative (randNEC) equivalence constraints, respectively. These randomly generated equivalence constraints were consistent with the task-assigned categories, but no attempt was made to control the information they provide as a group (i.e., their selection was random). In a sense, these random constraint conditions were designed to represent expected real-world scenarios in which the classifier is provided with haphazard constraints and not those that are necessarily most useful for good categorization.

Method

Participants

Participants in all three experiments were undergraduate or graduate students from the Institute of Life Sciences at the Hebrew University of Jerusalem. Participants were randomly assigned to the different experiments and experimental conditions, and did not participate in more than one Experiment. Twelve university students participated in the first experiment (mean age = 23.8, SD = 1.9), seven males and five females, with normal or corrected-to-normal vision.

Materials

Computer-generated pictures of “alien creature faces” were used as stimuli, as shown in Fig. 3. Each face was characterized by a unique combination of five potentially task-relevant dimensions: shape of chin, nose and ears, and color of skin and eyes. We designed 10 sets of 32 alien face stimuli such that for each set, all combinations of 5 binary dimensions were presented in each of the 10 experimental trials. All sets were used in each experimental condition in each one of the three experiments. Two or three (of the 5 possible) dimensions were selected as relevant for category definition on each trial, so that positively constrained pairs of objects had to have the same features (values) for all relevant dimensions and negatively constrained pairs had to differ in at least one of these. Stimuli were presented on a 22” high-resolution computer screen, using specially designed software that enabled both simultaneous presentation of many stimuli and the recording of participants’ reactions.

Procedure

All participants performed the three experimental conditions in a within-subject blocked-experiment design. Participants were told that during each experimental trial they would have to learn which of the 32 “alien creatures” (test stimuli) belonged to the same tribe as the one identified as “chief” (a standard representing the target category). They were instructed that each task (trial) in the

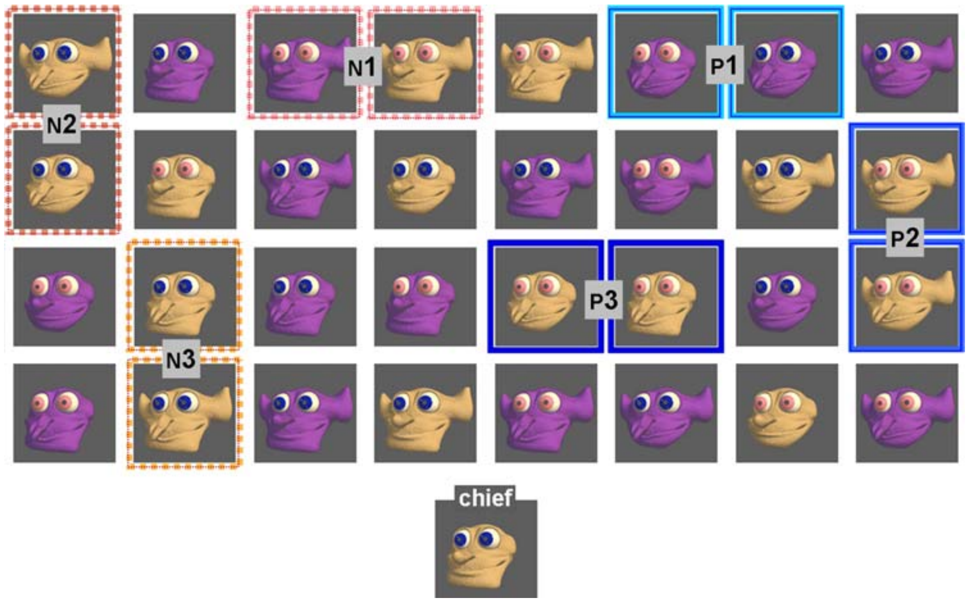

Fig. 3 Example of stimulus configuration on one specific trial. Participants decided which of the 32 test stimuli belong to the chief's tribe. Clues (constraints) were presented as frames surrounding pairs of exemplars. Positive and negative equivalence constraints (PECs and NECs) are illustrated, respectively, as *solid lines*, marked P1–P3, and *dashed lines*, marked N1–N3. Note that in the experiment, the two types of constraints never appeared together in the same trial. Highly informative constraints (Experiment 2 and 3), as illustrated here, present pairs of images that differ in only one feature. In the current example, participants had to learn that skin color and ear shape are relevant for categorization. Specifically, NEC N1 informs participants that skin color is a relevant dimension because it is the only dimension discriminating between the two exemplars. Similarly, N2 and N3 both imply that ear shape is relevant for categorization. P1, P2, and P3 inform participants that eye color, nose shape, and chin shape are not relevant for categorization since these features are different in pairs that belong to the same tribe. In the highly informative constraint task, as in the current example, all the information needed for proper categorization was provided (for either NECs, or PECs, separately; see text)

experiment was independent and would necessitate learning a new way of categorizing the aliens into tribes. Participants were not informed that for each trial two or three dimensions were chosen as trial-relevant. In general, we did *not* give subjects specific instructions which clarify the optimal categorization strategy or the structure of the categories; rather, participants were simply told that during each trial they will have to use the clues provided for identifying the chief tribe members. Participants were also instructed that they will have limited time to respond, and that they should perform the task not only accurately, but also as quickly as possible.

In general, clues (equivalence constraints) were provided as colored frames around pairs of aliens, indicating that the members of the pair belong to different tribes (randNEC condition) or the same tribe (randPEC condition). Figure 3 shows an example of an experimental trial. On each trial, three constraints appeared for 20 s together with the ensemble of alien faces. All the trial's constraints were presented simultaneously in order to allow participants to integrate the information provided by more than one constraint, without being affected by memory load. After 20 s the constraints were removed and the alien faces shuffled. Participants were then given 50 s to select (by drag-and-drop) those aliens that he or she thought belonged to the chief's tribe. The trial was then terminated and the next experimental trial began.

Even without using the information presented in the Equivalence Constraints, subjects could perform the categorization task by simply using an associative categorization strategy based on some idiosyncratic similarity measure. That is, for the chief's tribe they could choose those aliens that resembled the chief in some way. Therefore, we first tested participants on the "no equivalence constraints" (noEC) condition. In this condition participants performed the categorization task in a totally unsupervised manner: i.e., without being provided with either NECs or PECs. Performance in this condition was evaluated by tabulating the match between the tribe members selected by the participant and the expected tribe members according to the task pre-selected relevant dimensions.

After performing the noEC condition, participants performed the randomly selected NEC and PEC tasks (randNEC and randPEC, in counter-balanced order). In these experimental conditions the constraints were consistent with the computer-assigned alien creature categories. That is, there was no assignment of a NEC to two stimuli that belong to the same category or assignment of a PEC to two stimuli from two different categories. However, we made no attempt to select the three constraints in a way that maximized the

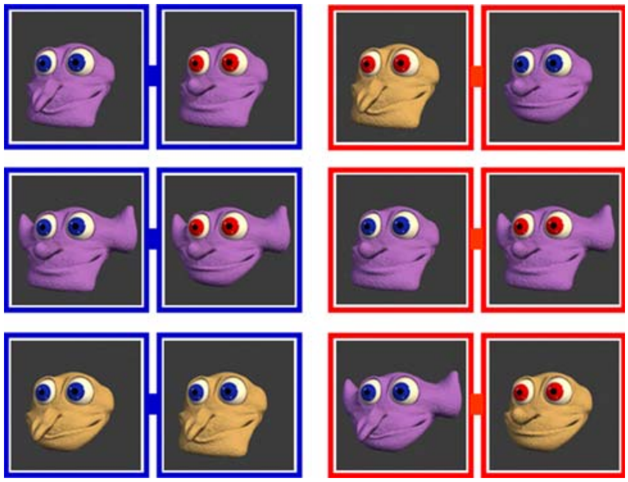

Fig. 4 Examples of randPECs (*left*) and randNECs (*right*). Randomly selected constraints are applied to the same experimental trial depicted in Fig. 3, but here the constraints are randomly selected. In this trial, participants had to identify ear shape and skin color as the relevant dimensions for categorizing the aliens. Thus, valid positively constrained pairs (PECs) include aliens with the same ear shape and skin color and with either identical or differing chin shape, nose shape and eye color. Generally, as few as three such pairs suffice to identify the irrelevance of the latter three dimensions and thus the relevance of the first two. On the other hand, valid negatively constrained pairs (NECs) will include aliens with either different ear shape or different skin color. These were usually non-informative since to be informative, the pair could not differ on any other dimension. As can be seen in these examples, the task-relevant dimensions can be easily identified from the three randPECs, but not from the randNECs, since the pairs differ also on non-relevant dimensions

information provided for optimal performance (identifying exactly all the trial-relevant dimensions). Note that for the reasons mentioned in the Introduction, in the randPEC condition the information provided by three randomly selected constraints almost always sufficed for identifying the task-relevant dimensions. This was not the case for randNECs, where the information provided was almost as poor as in the noEC condition. See Fig. 4 for examples of random PECs and NECs.

At the beginning of each experimental condition, participants performed an example trial in which they received a brief technical explanation about how they should perform the experiment and also about the identity of the constraints—whether the two constrained alien creatures are from the same tribe (PEC condition), or from different tribes (NEC condition).

Results and discussion

Performance measures

Participant performance is described by the Hit rate (the number of correctly selected "chief's tribe members" relative to the total number of "tribe members") and False-Alarm rate (the number of mistakenly selected "non-tribe members" relative to the total number of "non-tribe members"). Note that the average number of possible Hits in a trial is 6, and the average number of possible correct rejections is 26 (32 target stimuli in total). Therefore the False-Alarm values are expected to be relatively small compared to the Hit values.

To evaluate further participant sensitivity (i.e., their ability to discriminate between categories) we used the A' nonparametric sensitivity measure (Grier 1971; Stanislaw and Todorov 1999): A score of $A' = 0.5$ represents poor ability to discriminate between categories, whereas a score of $A' = 1$ represents perfect ability to discriminate between categories. Scores between 0 and 0.5 represent response confusion. Note that due to the differences in the prior probabilities of Hits versus False-Alarms, in the current experimental tasks chance performance is expected to result in A' higher than 0.5. A' is calculated as follows:

$$A' = 0.5 + \left[\operatorname{sign}(H - F) \times \frac{(H - F)^2 + |H - F|}{4 \times \max(H, F) - 4 \times H \times F}\right]$$

where H denotes Hit rate, and F denotes False-Alarm rate. Participant reaction time was also recorded. Reaction time represents the average time it took for participants to detect and select each member in the target category, starting from the point when the constraints were removed and the stimuli were shuffled.

Results

An ANOVA revealed a significant effect of the experimental condition both on participants' Hits $F(2, 11) = 11.84$, $p < 0.001$, $\eta_p^2 = 0.52$, and False Alarms, $F(2, 11) = 4.67$, $p < 0.05$, $\eta_p^2 = 0.30$. Post-Hoc analysis using within-subject t test showed that randomly chosen positive constraints serve to improve performance, while randomly chosen negative constraints do not contribute any more to participant learning than does the condition with no constraints at all. There was a significantly higher Hit rate in the randPEC condition ($M = 0.57$, $SD = 0.15$) compared with the noEC condition ($M = 0.39$, $SD = 0.10$), $t(11) = 4.12$, $p < 0.005$, $d = 2.48$, as well as compared with the randNEC condition ($M = 0.44$, $SD = 0.14$), $t(11) = 3.54$, $p < 0.005$, $d = 2.13$. Similarly, the False-Alarm rate in the randPEC condition ($M = 0.10$, $SD = 0.05$) was significantly lower than in the noEC condition ($M = 0.14$, $SD = 0.04$), $t(11) = 2.64$, $p < 0.05$, $d = 1.59$, or in the randNEC condition ($M = 0.15$, $SD = 0.08$), $t(11) = 2.52$, $p < 0.05$, $d = 1.52$. On the other hand, there was no significant difference between the randNEC and noEC conditions in either the Hit rate $t(11) = 1.40$, $p = 0.19$ or False-Alarm rate $t(11) = 0.56$, $p = 0.59$.

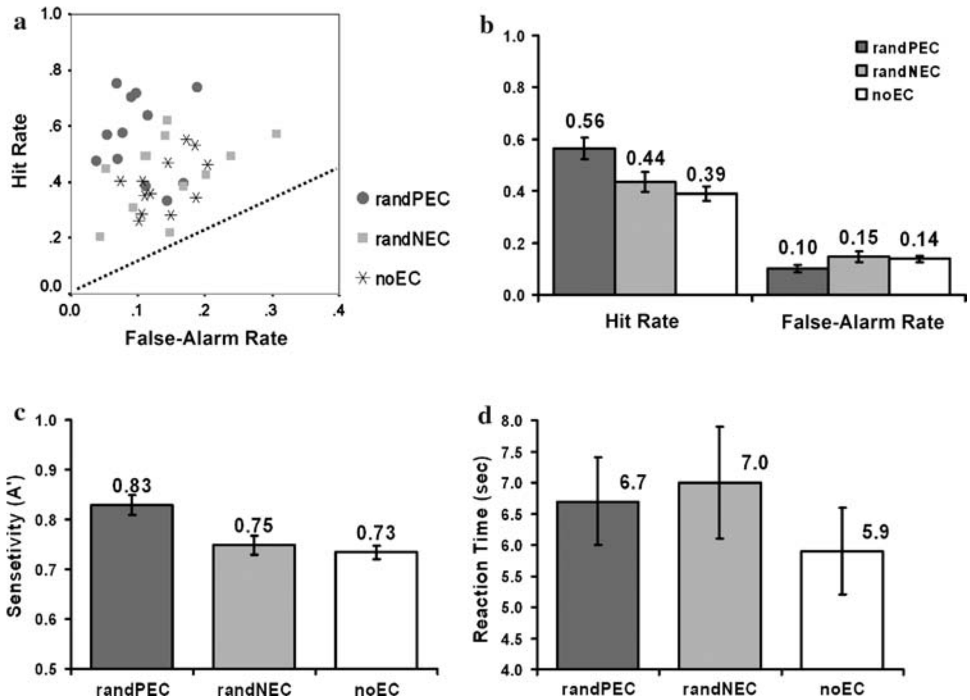

Fig. 5 Experiment 1. Performance without equivalence constraints (noEC) or with randomly chosen positive or negative equivalence constraints (randPEC or randNEC). **a** The receiver operating characteristics (ROC) diagram, plotting Hit rate (ordinate) versus False Alarm rate (abscissa). Each point represents one participant's performance in the specified experimental condition. Distance from the dashed line represents participant sensitivity, with points near the line representing random stimuli selection when assuming identical probability for Hits and FAs. (Note that the abscissa is limited to the range 0–0.4, since, as expected, there were relatively few FAs; see text.) **b** Mean Hit and False-Alarm rates in the three experimental conditions. **c** Mean sensitivity (A'). **d** Mean reaction time (in seconds). *Error bars* in all figures are standard errors of the mean

This effect of random constraints type was also apparent when evaluating participants' sensitivity using A' (defined above). An ANOVA showed a significant difference between conditions, $F(2, 11) = 16.27$, $p < 0.001$, $\eta_p^2 = 0.60$. Post-Hoc analysis using paired sample t tests revealed that sensitivity in the randPEC condition ($M = 0.83$, $SD = 0.02$) was significantly higher than in either the randNEC ($M = 0.75$, $SD = 0.02$), $t(11) = 4.81$, $p < 0.001$, $d = 2.90$, or noEC conditions ($M = 0.73$, $SD = 0.01$), $t(11) = 4.33$, $p < 0.005$, $d = 2.61$. There was no significant difference between sensitivity in the randNEC and noEC conditions, $t(11) = 1.02$, $p = 0.33$. Nevertheless, there was no significant difference in reaction time between the three conditions (ANOVA: $F(2, 11) = 2.24$, $p = 0.13$). These results are illustrated in Fig. 5.

Discussion

The goal of Experiment 1 was to measure baseline performance. Participant performance in the noEC condition represents the expected categorization performance when simply using an idiosyncratic associative categorization strategy. In this way we can estimate the contribution of the information provided by equivalence constraints in the other conditions more appropriately. The results confirmed the theoretical conclusion that a set of random PECs is more informative than a set of random NECs. This finding may explain the results of other studies in which randomly chosen PECs lead to better performance than do randomly chosen NECs; e.g. (Whitman and Garner 1962 used sequences of same-category vs. alternating-category exemplars and found that the former leads to better performance). Our results also confirm the expectation that a small number of randNECs are poorly informative. The absence of significant differences in reaction time suggests that when provided with random PECs, participants can perform the categorization task much more accurately, but also nearly as quickly, as when operating with an unconstrained idiosyncratic associative categorization strategy, as they were left to do in the noEC condition.

Experiment 2: highly informative sets of equivalence constraints

In the second experiment participants performed categorization tasks similar to those in Experiment 1, but in this experiment both the PECs and the NECs were deliberately selected so as to provide all the information needed for perfect performance. We call the two types of constraints used in this experiment highly informative PECs and NECs (highPEC and highNEC conditions). The two types of constraints were provided to the participants separately in these two experimental conditions. The goal here was to determine participant inherent proficiencies in the use of PECs and NECs. In this experiment we used a between-subject design to ensure that experience with one type of constraints would not influence performance with the other.

Method

Participants

Eighty university students participated in the experiment (mean age = 24.2, SD = 2.8), 32 males and 48 females, with normal or corrected-to-normal vision. Participants were randomly assigned to the two experimental groups (highPEC or highNEC), in a between-subject design. The large sample in this experiment was essential since the statistical analysis used here included not only simple mean comparison, but also higher order analyses of homogeneity of variance and normality tests. To ensure the reliability of such analysis, large samples are required.

Materials

Identical to Experiment 1.

Procedure

The procedure in Experiment 2 was similar to the procedure in Experiment 1 except for the nature of the PECs and NECs that were provided. In this experiment PECs and NECs were deliberately selected so that each constraint would identify only one dimension as irrelevant (in the case of a PEC) or as relevant (in the case of a NEC). A "highly informative PEC" (highPEC) is composed of a pair of "aliens" from the same category (tribe) that differ in only one irrelevant dimension (e.g. the shape of their noses), so that the constraint enables participants to identify that this differentiating dimension is irrelevant for categorization due to the within-category variation in this dimension. A "highly informative NEC" (highNEC) is composed of a pair of aliens from two different categories (tribes) such that the pair of aliens differ in only one dimension, which should be identified as a relevant dimension due to the between category variation in the dimension (the only dimension enabling the discrimination between two stimuli from different categories). In this experiment participants could identify all the trial-relevant dimensions by integrating the information from the highPECs or highNECs provided, and therefore they could (in principle) perform the categorization task perfectly in both conditions. See Fig. 3 for examples of highly informative NECs and PECs. In each trial, the pre-selected relevant dimensions were identical to those of the respective trial in Experiment 1.

Results and discussion

Performance measures

Identical to Experiment 1.

Results

Between subject t tests showed no significant differences between the Hit rate in the highPEC condition ($M = 0.64$, SD = 0.17) and the highNEC condition ($M = 0.57$, SD = 0.25), $t(78) = 1.49$, $p = 0.14$. On the other hand, the False-Alarm rate in the highPEC condition ($M = 0.12$, SD = 0.06) was significantly higher than in the highNEC condition ($M = 0.07$, SD = 0.06), $t(78) = 3.48$, $p < 0.001$, $d = 0.79$ (see also Fig. 6). Nevertheless, participant sensitivity (A') in the highPEC condition ($M = 0.85$, SD = 0.07) was not significantly different than in the highNEC condition ($M = 0.83$, SD = 0.13), $t(78) = 0.85$, suggesting that the differences in the False-Alarm rates between the two conditions did not derive from a higher sensitivity in the highNEC group, but rather mainly from differences in response bias, where participants in the highPEC condition had a greater tendency to produce more False-Alarms together with a few more Hits (although the difference in the Hit rate was not significant) so that their categorization strategy can be described as more liberal. Later we will address this difference in more detail. These results are illustrated in Fig. 6a–c.

Performance in the two experimental conditions differed in participant reaction time with RT in the highPEC condition ($M = 6.9$ s, SD = 2.5 s) significantly shorter than in the highNEC condition ($M = 8.8$ s, SD = 4.1 s), $t(78) = 2.54$, $p < 0.05$, $d = 0.58$. Generally, there were no differences in performance when comparing experimental trials with two relevant dimensions with those with three relevant dimensions—except that the expected main effects showed poorer performance in trials with three relevant dimensions, which may be perceived as more difficult. The only exception was a significant interaction between the highPEC and highNEC experimental conditions (between-subject variable) and the number of relevant dimensions (within-subject variable), $F(1, 77) = 24.58$, $p < 0.001$, $\eta_p^2 = 0.24$. Post-hoc t tests revealed that while there was no significant difference in reaction time between trials with two relevant dimensions ($M = 6.7$ s, SD = 3.1 s) and

Fig. 6 Experiment 2. Performance with highly informative positive or negative equivalence constraints (highPEC or highNEC). **a** The receiver operating characteristics diagram showing largely overlapping results for highlyinformative PECs and NECs. **b** Mean hit and false-alarm rates. **c** Mean sensitivity (A'). **d** Mean reaction time (in seconds)

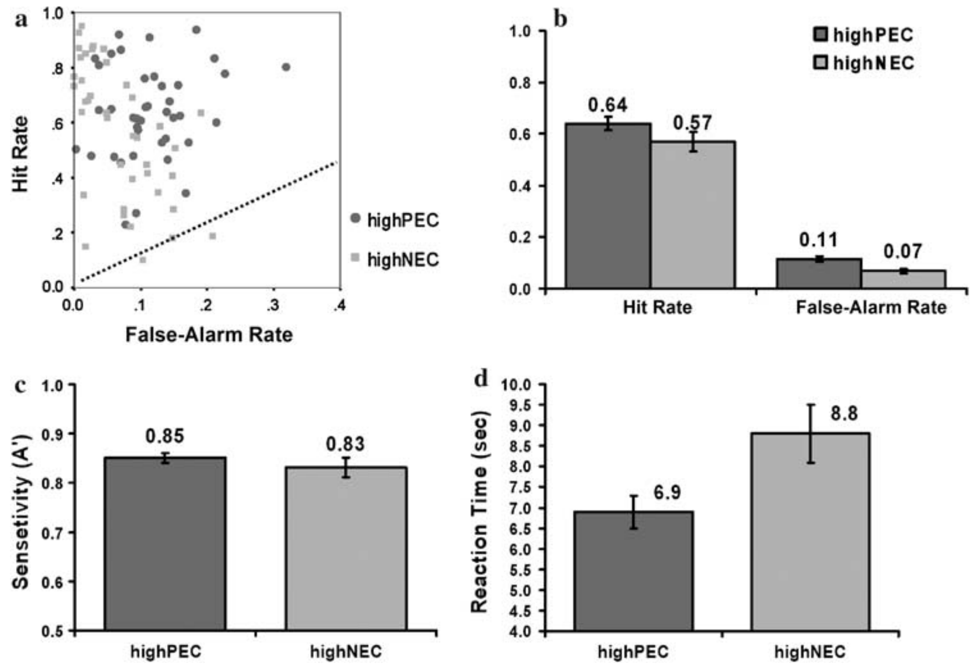

trials with three relevant dimensions ($M = 7.0$ s, $\mathrm{SD} = 2.7$ s) in the highPEC condition, in the highNEC condition reaction time in trials with two relevant dimensions ($M = 7.0$ s, $\mathrm{SD} = 3.8$ s) was significantly shorter than reaction time in trials with three relevant dimensions ($M = 10.6$ s, $\mathrm{SD} = 5.2$ s). These results are illustrated in Fig. 6d.

More importantly, the highPEC and highNEC groups also significantly differed in the distribution of their Hit rates. Levene's test for homogeneity of variances showed that the Hit rate in the highNEC condition is more variable across participants compared to the highPEC condition, $F(78) = 8.93$, $p < 0.005$. This difference is also apparent in the A' standard-deviation, with a smaller standard-deviation in the highPEC condition than in the highNEC condition, $F(78) = 13.94$, $p < 0.001$. The Shapiro–Wilk test of normality further shows that although in the highPEC condition, sensitivity is normally distributed, $W(40) = 0.95$, $p = 0.11$, the distribution of sensitivity in the highNEC condition differs significantly from normal, $W(40) = 0.89$, $p < 0.001$. As can be seen in Fig. 7a, while in the highPEC condition the sensitivity distribution shows good fit with the expected normal curve and most participants show good sensitivity, in the highNEC condition there is a poor match with the expected normal.

This divergence from the expected normal distribution is also illustrated in Fig. 7b, where we plot on top of each ROC diagram a horizontal line representing the median Hit rate and a vertical line representing the median False-Alarm rate. It is clearly seen that participants in the highNEC group (Fig. 7b-right) are for the most part separated into two distinct subgroups: participants with poor performance (lower right quadrant) versus those with good performance (upper left quadrant). This is not the case in the highPEC condition (Fig. 7b-left) in which performance is distributed evenly around and relatively close to the crossing point of the medians. Thus, there is an important difference between the use of PECs and NECs: While most participants correctly used highPECs in the category learning tasks, performance in the highNEC condition varied—with about half of the participants succeeding in proper use of these highly informative NECs, even surpassing the performance of the highPEC group, and the others failing to derive any benefit from these highNECs.

By comparing the results of Experiment 2 with highly informative equivalence constraints, to those of Experiment 1 with randomly selected constraints, we find that in the highNEC condition performance was significantly better than in the randNEC condition, while there was no significant difference between the randPEC and highPEC conditions. The superior performance in the highNEC condition stems from both more Hits and fewer FAs than in the randNEC condition, as follows: The Hit rate in the highNEC condition ($M = 0.57$, $\mathrm{SD} = 0.25$) was significantly higher than in the randNEC condition ($M = 0.44$, $\mathrm{SD} = 0.14$), $t(50) = 2.44$, $p < 0.05$, $d = 0.84$. Similarly, the False-Alarm rate in the highNEC condition ($M = 0.07$, $\mathrm{SD} = 0.06$) was significantly lower than in the randNEC condition ($M = 0.15$, $\mathrm{SD} = 0.08$), $t(50) = 3.88$, $p < 0.001$, $d = 1.10$. These differences in the Hit and False-Alarm rates were also apparent when comparing participant sensitivity in the two cases: In the

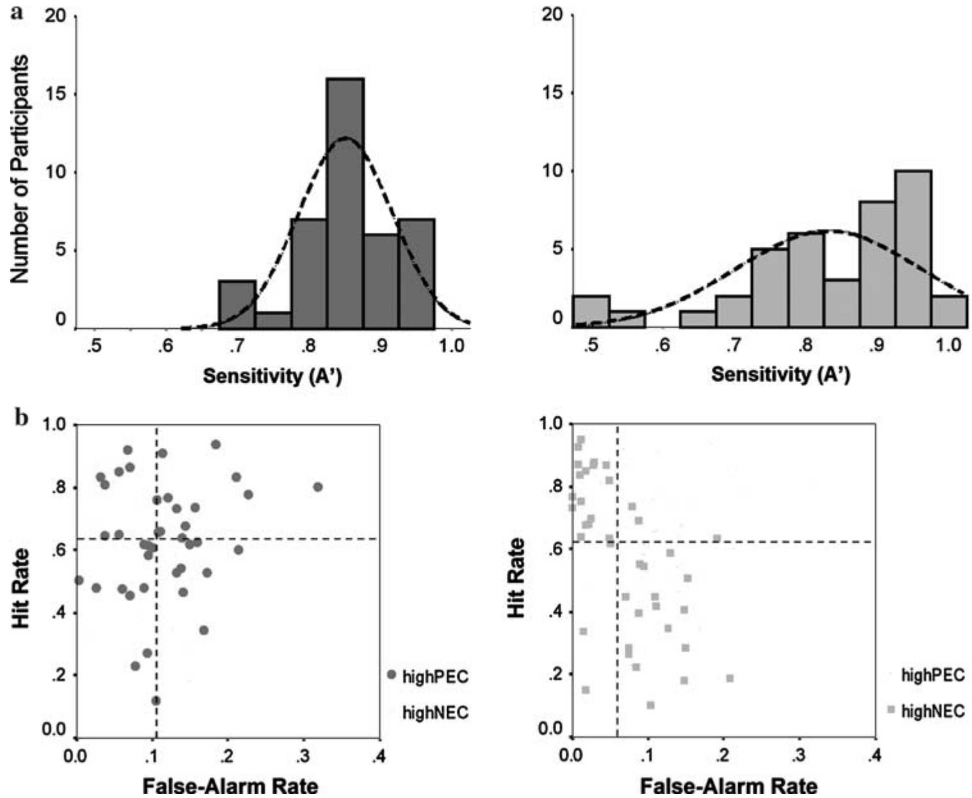

Fig. 7 **a** Sensitivity distribution in the highPEC (*left*) and highNEC (*right*) conditions of Experiment 2. The *horizontal axis* represents participant sensitivity (A') and the vertical axis represents the number of participants. *Dashed curves* represent the expected normal curves calculated from each group mean and standard deviation. **b** Receiver operating characteristic diagrams for the highPEC (*left*) and highNEC (*right*) conditions. *Dashed lines* represent the median Hit (*horizontal lines*) and False-Alarm (*vertical lines*) rates in each condition

Table 1 Summary comparing the highly informative and random constraint conditions for positive and negative constraints (significant values are indicated by *)

	PEC		NEC	
	Mean ± SD	*t* test	Mean ± SD	*t* test
Hits				
High	0.64 ± 0.17	$t(50) = 1.38$	0.57 ± 0.25	$t(50) = 2.44$
Rand	0.57 ± 0.15	$p = 0.17$	0.44 ± 0.14	$p^* < 0.05$
False-alarms				
High	0.12 ± 0.06	$t(50) = 0.69$	0.07 ± 0.06	$t(50) = 3.88$
Rand	0.10 ± 0.05	$p = 0.50$	0.15 ± 0.08	$p^* < 0.001$
Sensitivity (A')				
High	0.85 ± 0.07	$t(50) = 1.00$	0.83 ± 0.13	$t(50) = 2.13$
Rand	0.83 ± 0.07	$p = 0.32$	0.75 ± 0.02	$p^* < 0.05$
Reaction time				
High	6.9 ± 2.5 s	$t(50) = 0.48$	8.8 ± 4.1 s	$t(50) = 1.44$
Rand	6.7 ± 3.2 s	$p = 0.63$	7.0 ± 3.1 s	$p = 0.16$

Note that there were no significant differences between the two PEC conditions, while performance with highly informative NECs is significantly better than with randomly selected NECs

highNEC condition ($M = 0.83$, SD = 0.13) sensitivity was significantly higher than in the randNEC condition ($M = 0.75$, SD = 0.02), $t(50) = 2.13$, $p < 0.05$, $d = 0.60$. Taken together, these findings confirm that in the highNEC condition constraints provide more information than in the randNEC condition, and that, in general, participants successfully used this information that enabled better performances. Performance in the highPEC condition did not differ significantly from that in the randPEC condition (see Table 1), suggesting that a deliberate selection of an informative set of PECs is not more beneficial than a randomly selected set of PECs.

Summary

After showing in Experiment 1 that sets of random NECs are less informative than sets of random PECs, we designed Experiment 2 to test whether this differentiating property of equivalence constraints affects the way people perceive and integrate PECs and NECs in general. It is possible that experience with natural conditions where NECs are generally not informative may lead to a lack of expertise in the use of NECs, and a resulting inability to use even highly informative NECs. Alternatively, despite their inexperience with informative NECs, participants may be sufficiently skilled and flexible so that they will be able to extract the information supplied to them when NECs are highly informative. In fact, as we show below, NECs may actually be easier to use than PECs. Furthermore, the very lack of experience may allow participants to use NECs in a more innovative and informative fashion than PECs.

The results in the NEC condition clearly divide our participants into two groups, with one group lacking the ability to use highly informative NECs efficiently, and the other succeeding brilliantly in their use—surpassing even the performance with PECs. Such variability is not observed in the PEC condition. These results may suggest that NECs and PECs are used quite differently: In the PEC condition, the unimodal sensitivity distribution with its relatively small standard deviation, together with the relatively fast reaction time, provides solid evidence that category learning from PECs is done intuitively by most people. In contrast, in the NEC condition, the somewhat bimodal distribution and relatively large standard deviation, together with the long reaction time that was also highly dependent on task difficulty (two vs. three relevant dimensions), indicate that category learning from even highly informative NECs is not naturally performed and requires expertise that only some people have. This ability to correctly use highNECs for category learning tasks results in nearly perfect performance.

Experiment 3: highly informative equivalence constraints with directions

We concluded from the results of Experiment 2 that participants may have different abilities for reasoning about informative NECs—perhaps due to this type of constraint being rare in natural conditions. If this is so, then guiding people in the use of highly informative NECs may improve performance. This is the goal of Experiment 3. In this experiment, participants performed a categorization task identical to the one that was performed in Experiment 2, using exactly the same sets of highly informative PECs and NECs. The only difference between the two experiments was that in the current experiment we also provided participants with "meta-knowledge"—explicit directions for a categorization strategy enabling perfect performance. If the difference between the two highNEC subgroups in Experiment 2 was due to the fact that some participants did not know how to use these constraints, then giving them directions for their use should bring performance of all participants to the level of the better subgroup. In addition, Experiment 3 may help evaluate the findings of Experiment 2 with regard to the use of PECs. More specifically, we wanted to know whether the pattern of performance of participants in the highPEC condition in Experiment 2 truly represents the expected performance when in possession of the optimal rule-based categorization strategy.

Method

Participants

Twelve university students participated in the experiment; mean age = 23.9, SD = 5.4, 7 males and 5 females, with normal or corrected-to-normal vision.

Materials

Identical to Experiments 1 and 2.

Procedure

The procedure in Experiment 3 was identical to that of Experiment 2 with exactly the same sets of highPECs and highNECs. The only difference between the two experiments was that in the instructions provided during the example trial of each condition, participants in Experiment 3 were also directed how they should integrate the information provided by the equivalence constraints. The directions were straightforward and simple, and all participants easily learned the principles provided. More specifically, before performing the highPEC condition, participants were informed that they should exclude the dimension discriminating between each two constrained exemplars, since this dimension was necessarily irrelevant for the categorization task, and reserve judgment about the rest of the dimensions, with identical features, since they may or may not be relevant. Before performing the highNEC condition, participants were informed that they should take into account the dimension discriminating between each two constrained exemplars because, as the only differentiating dimension it must be relevant for the categorization task. Participants performed the experiment as a within-subject experimental design with the order of the two experimental conditions being counter-balanced.

Results and summary

Performance measures

Identical to Experiment 1 and 2.

Results

Surprisingly, performance in the directed highNEC condition was superior to the performance in the directed highPEC condition, as shown in Fig. 8. That is, when directions were given, the usefulness of the informative positive constraints was not improved, and the information provided by the negative constraints not only improved performance with these constraints, but such performance also surpassed that with the positive constraints. Specifically, in contrast to Experiment 2, the Hit rate in the directed negative constraint condition ($M = 0.86$, SD $= 0.10$) was significantly higher than with positive constraints ($M = 0.68$, SD $= 0.18$), $t(11) = 3.09$, $p < 0.05$, $d = 1.86$. Similarly, the False-Alarm rate in the directed-highNEC condition ($M = 0.03$, SD $= 0.04$) was significantly lower than with positive constraints ($M = 0.08$, SD $= 0.05$), $t(11) = 3.07$, $p < 0.05$,

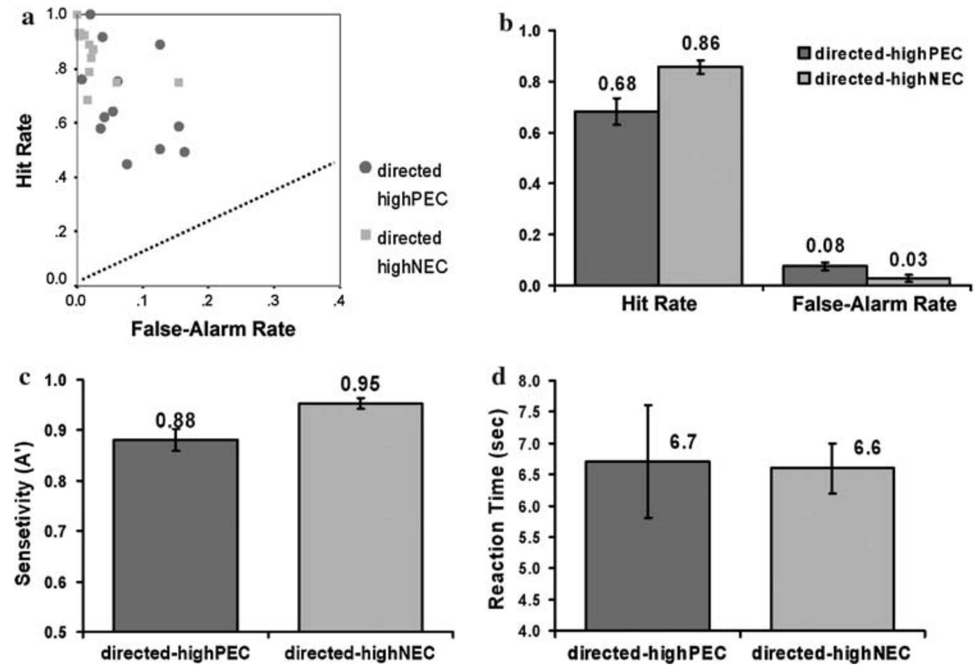

Fig. 8 Experiment 3. Performance following directions for optimal use of highly informative positive or negative equivalence constraints (directed-highPEC or directed-highNEC conditions). **a** The receiver operating characteristics diagram. **b** mean hit and false-alarm rates. **c** Mean sensitivity (A'). **d** Mean reaction time (in seconds)

$d = 1.85$. As a consequence, sensitivity in the directed-highNEC condition ($M = 0.95$, SD $= 0.04$) was also higher than in the directed-highPEC condition ($M = 0.88$, SD $= .07$), $t(11) = 3.29$, $p < 0.01$, $d = 1.98$. This superior performance did not occur at the cost of slower response, as there was no significant difference in reaction time between the directed-highNEC condition ($M = 6.6$ s, SD $= 1.5$ s) and directed-highPEC condition ($M = 6.7$ s, SD $= 3.2$ s), $t(11) = 0.10$ (see also Fig. 8d).

We now compare the non-directed highly informative equivalence constraint conditions of Experiment 2 with the directed highly informative equivalence constraint conditions of Experiment 3. Between-subject t-tests revealed that providing participants with directions affected mostly the way highNECs were used in categorization tasks but had almost no effect on the way highPECs were used for such tasks. More specifically, the Hit rate in the directed-highNEC condition ($M = 0.86$, SD $= 0.10$) was significantly higher than without directions ($M = 0.57$, SD $= 0.25$), $t(50) = 3.94$, $p < 0.001$, $d = 1.12$. The False-Alarm rate in the non-directed highNEC condition ($M = 0.07$, SD $= 0.06$) was significantly higher than in the directed condition ($M = 0.03$, SD $= 0.04$), $t(50) = 2.65$, $p < 0.05$, $d = 0.75$. Sensitivity in the directed-highNEC condition ($M = 0.95$, SD $= 0.04$) was also significantly higher than in the non-directed condition ($M = 0.83$, SD $= 0.13$), $t(50) = 5.30$, $p < 0.001$, $d = 1.50$. This improvement in categorization accuracy did not occur at the cost of longer reaction time. In fact, reaction time was shorter in the directed-highNEC condition ($M = 6.6$ s, SD $= 1.5$ s) than in the non-directed highNEC condition ($M = 8.8$ s, SD $= 4.1$ s), $t(50) = 2.77$, $p < 0.01$, $d = 0.78$. In comparison to this across the board improvement with directions in the highNEC condition, there was no significant improvement when participants were provided with directions together with highPECs. Table 2 summarizes the impact of providing directions by comparing the results of Experiments 2 and 3.

Table 2 Summary comparing the directed and non-directed highly informative constraint conditions for positive and negative constraints (significant values are indicated by *)

	HighPEC		HighNEC	
	Mean ± SD	t test	Mean ± SD	t test
Hits				
Non-directed	0.64 ± 0.17	$t(50) = 0.75$	0.57 ± 0.25	$t(50) = 3.94$
Directed	0.68 ± 0.18	$p = 0.45$	0.86 ± 0.10	$p^* < 0.001$
False-alarms				
Non-directed	0.12 ± 0.06	$t(50) = 1.95$	0.07 ± 0.06	$t(50) = 2.65$
Directed	0.08 ± 0.05	$p = 0.06$	0.03 ± 0.04	$p^* < 0.05$
Sensitivity (A')				
Non-directed	0.85 ± 0.07	$t(50) = 1.36$	0.83 ± 0.13	$t(50) = 5.30$
Directed	0.88 ± 0.07	$p = 0.18$	0.95 ± 0.04	$p^* < 0.001$
Reaction time				
Non-directed	6.9 ± 2.5 s	$t(50) = 0.17$	8.8 ± 4.1 s	$t(50) = 2.77$
Directed	6.7 ± 3.2 s	$p = 0.87$	6.6 ± 1.5 s	$p^* < 0.01$

Note that there were no significant differences between the two PEC conditions, while directions provided with highly informative NEC constraints significantly improved performance

Discussion

Providing highNECs together with directions for their use is extremely helpful in boosting accuracy and response time. The impact of directions is manifested not only in the improved performance in the directed-highNEC condition compared to the highNEC condition, but also in the relatively much more homogeneous performance in the directed-highNEC condition. In contrast, except for a moderate and barely significant reduction in the False-Alarm rate, performance in the directed-highPEC condition was not significantly improved compared to the non-directed highPEC condition. Performance in these two conditions is also similarly homogeneous.

Further comparisons of experiments 1–3

In order to investigate further the observed non-homogeneous performance in the highNEC condition of Experiment 2, we divided the highNEC group into two subgroups of 20 participants each—the highNEC-poor (participants with relatively low performance) and the highNEC-good (participants with relatively high performance), separated by the median sensitivity ($A' = 0.86$) of the highNEC group (see Experiment 2, Results). It is important to stress that this separation into two groups is artificial, and the A' value of 0.86 does not necessarily represent an objective borderline separating poor performers from the good ones. Nevertheless, using a large sample insures that this observed median is a good approximation for the expected median performance in the population, (taking into account the type of population from which the participants were sampled).

In order to understand better the source of this apparently bimodal performance, and the resulting division into two subgroups, we compared the separate performance of these two subgroups with those of participants who were given practically non-informative constraints, on the one hand, and with participants who were given the best possible information, (including both highly informative constraints and directions for their use), on the other. In other words, we compared the performance of the highNEC-poor and highNEC-good subgroups (Experiment 2) to that of participants in the randNEC (Experiment 1) and directed-highNEC (Experiment 3) conditions. In Fig. 9a, we replot the randNEC points of Fig. 5a and the directed-highNEC points of Fig. 8a. We also reproduce in this graph the median dividing lines between the highNEC-good and highNEC-poor performers of Fig. 7b-right. Clearly, the randNEC points fall neatly within the lower-right quadrant—where the data of the highNEC-poor performers are situated (see Fig. 7b-right) and the directed-highNEC points fall neatly in the upper-left quadrant, the location of the data of the highNEC-good performers. Similarly, the sensitivity of the highNEC-good subgroup is similar to that of the directed-highNEC group, and the sensitivity of the highNEC-poor subgroup matches that of the randNEC group, as shown in Fig. 9b.

Specifically, sensitivity in the highNEC-good subgroup ($M = 0.93$, $SD = 0.03$) was as high as in the directed-

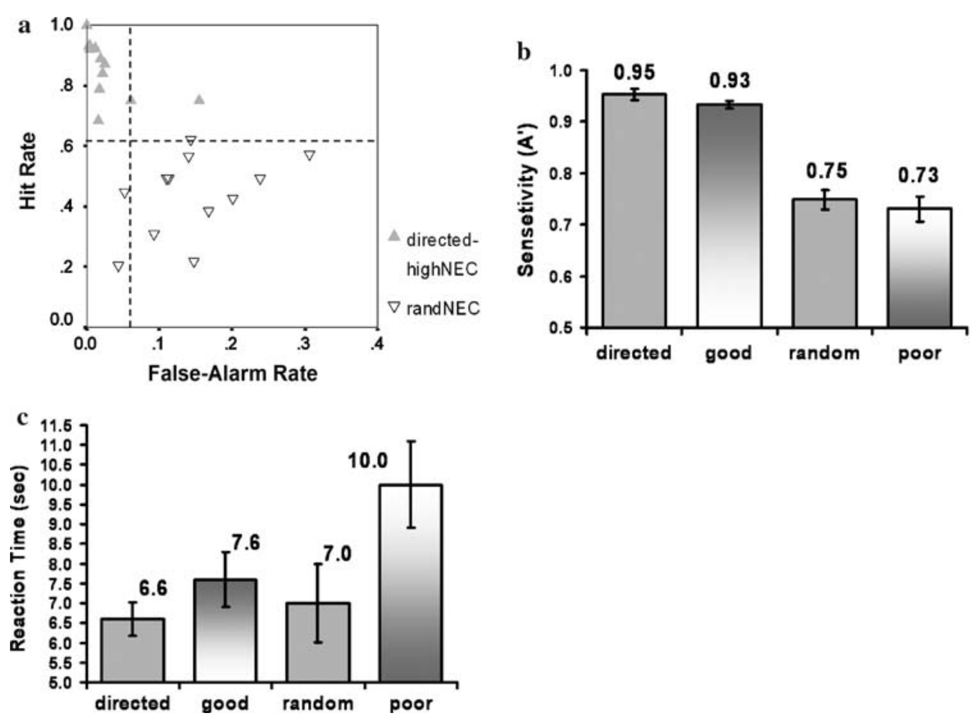

Fig. 9 Between-experiment comparisons. **a** The receiver operating characteristics diagram of the directed-highNEC (Exp. 3, Fig. 8a) and randNEC (Exp. 1, Fig. 5a) conditions. *Dashed lines* represent median Hits (*horizontal line*) and False-alarms (*vertical line*) as they were calculated for the highNEC condition (see Exp. 2, Fig. 7b). Note the clear separation of these results into upper-left and lower-right quadrants, respectively. **b** Mean sensitivity (A') for participants receiving NECs in each of the three experiments. **c** Mean reaction time (in seconds) for these groups of subjects

highNEC condition ($M = 0.95$, SD $= 0.04$), $t(30) = 1.70$, $p = 0.10$. At the same time, sensitivity in the highNEC-poor subgroup ($M = 0.73$, SD $= 0.11$) was as low as in the randNEC condition ($M = 0.75$, SD $= 0.07$), $t(30) = 0.52$. Also, reaction time in the highNEC-good subgroup ($M = 7.6$ s, SD $= 2.5$ s) was as fast as in the directed-highNEC condition ($M = 6.6$ s, SD $= 1.5$ s), $t(30) = 1.16$, $p = 0.25$, see Fig. 9c. On the other hand, the mean reaction time in the highNEC-poor subgroup ($M = 10.0$ s, SD $= 5.0$ s) was not as fast as in the randNEC condition ($M = 7.0$ s, SD $= 3.1$ s), but this difference was not highly significant, $t(30) = 1.91$, $p = 0.07$, as a result of the large variability in participants' reaction time in these two groups.

Discussion

Participants in the highNEC-good subgroup of Experiment 2 apparently implemented a similar or similarly effective categorization strategy as that used by participants in the directed-highNEC condition of Experiment 3. On the other hand, participants in the highNEC-poor subgroup from Experiment 2 failed to implement a useful categorization strategy, and they performed the categorization task just as the participants in the randNEC condition of Experiment 1, who received random constraints with low information value. The only difference was in reaction time, which was somewhat longer in the highNEC-poor subgroup than in the randNEC condition. This suggests that although participants in the highNEC-poor subgroup failed to properly use the information provided, they may have invested time trying to do it ineffectively.

General discussion

In the introduction we described inherent ecological differences between positive equivalence constraints (PECs) and negative equivalence constraints (NECs). Our main observation was that PECs are more informative than NECs. We then hypothesized that this fact may affect the way people process PECs and NECs in general. That is, the statistical difference in usability of NECs and PECs may lead people to expect (inherently and presumably unconsciously) the NECs not to be informative. This expectation may result in their superior use of PECs and thus their inability to process even informative NECs. The current research findings strongly support this hypothesis.

In Experiment 1, which was designed to evaluate baseline performance, we saw a clear advantage for category learning from randomly selected PECs compared to randomly selected NECs. Moreover, as expected from the theoretical background, random NECs were found to be poorly informative, enabling only categorization performance similar to that observed when participants merely performed associative categorization, as in the control condition without constraints.

Experiment 2 investigated whether the fact that PECs and NECs are differently informative affects the way people process these constraints when they are equally and highly informative. Results showed that deliberately selected PECs, containing all the information needed for perfect performance, were in fact not more beneficial than randomly selected PECs (which are also likely to contain all the information needed for perfecting performance). In contrast, deliberately selected informative NECs enabled much better performance than randomly selected NECs. Taken as a group, participants in the highNEC condition had similar sensitivities to those in the highPEC condition. The main differences were that the highNEC group had a slower mean reaction time, a lower False-Alarm rate, and an evidently but non-significantly lower Hit rate. It seems like highNECs lead participants to use a more conservative decision criterion at the cost of a longer reaction time.

Further analysis revealed an interesting dichotomy in the highNEC group: While in the highPEC group, sensitivity was normally distributed with a relatively small standard-deviation, in the highNEC condition, the sensitivity distribution was not unimodal and it had a relatively large standard-deviation. This pattern of performance in the highNEC condition was also apparent in the Hit and False-Alarm distribution patterns, showing that about half of the participants in the highNEC condition performed almost perfectly while the other half performed poorly, as though they had not received any informative constraints at all. In contrast, in the highPEC condition, both nearly perfect and poor performances were relatively rare. Instead, most participants showed reasonably good performance. Further testing of reaction time also revealed that in the highNEC condition, responses were not only slower than with PECs, but they were also highly dependent on task difficulty and individual participant sensitivity; namely, participants with high sensitivity also had faster reaction times. These findings clearly demonstrate that while the use of PECs is accomplished relatively easily and intuitively, many people have difficulty in using highNECs in category learning tasks.

Experiment 3 provided a number of surprising results. First of all, we found that the strategy for using highNECs could be readily learned via simple instructions, leading participants to nearly perfect performance. This change—compared to the non-directed highNEC case in Experiment 2—was probably due to improved performance of the potentially poor performing subgroup, bringing them up to the level of good performers. This result suggests that the failure of the poor performance subgroup in using

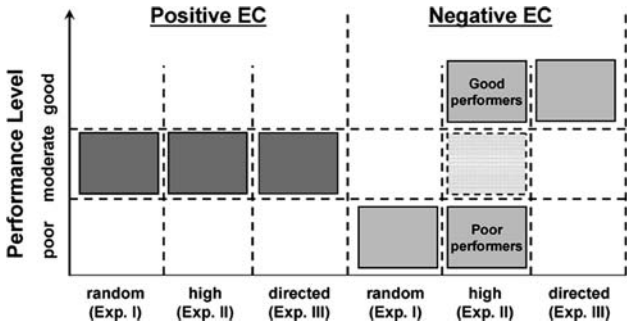

Fig. 10 Schematic summary of performance in the three experiments of this study—with randomly chosen constraints (I) or highly informative constraints, without (II) or with (III) directions for their use. The level of performance with Positive Equivalence Constraints (left) was similar and moderate in all of the three experiments. Deliberately selecting constraints for maximizing information (high-PEC, Exp. II) and providing participants with directions for how to use these constraints (directed-highPEC, Exp. III) did not improve performance compared to the randPEC condition (Exp. I). The pattern of performance with negative constraints (right) was different: While performance with randomly chosen constraints was poor (randNEC, Exp. I), deliberate selection of informative NECs (Exp. II) resulted in a bimodal distribution of performance with some participants performing poorly and others almost perfectly. Providing directions (Exp. III) resulted in near-perfect performance for most participants

highNECs was due to their inability to autonomously find the correct strategy, and not their inability to adopt new strategies. Still, it is surprising that a strategy for using highNECs was easily learned when instructions were provided, but many people (university students!) failed in intuitively implementing this strategy when performing the task without instruction.

Second, we found that giving similar instructions for the best strategy for using PECs did not improve performance and participants remained at quite good, but not perfect performance levels. This difference between the benefit of instructions for using PECs and NECs was rather unexpected, and supports our main conclusion that people use PECs, but not NECs, intuitively. These findings also help in rejecting the possibility that the pattern of performance observed in the highNEC group in Experiment 2 could have resulted from some confusion of the poor performers concerning the experimental setting. Figure 10 summarizes participant performance in the three experiments. Note that performance with PECs is similar in all three experiments, while with NECs, performance improves for some when we gave informative NECs, and for all when directions were added.

Strategies for using positive equivalence constraints

The lack of change when provided with instructions for using PECs (in Exp. 3) may be accounted for by one of the following: (1) participants' default strategy was similar to the rule-based strategy suggested by the directions, and so the "tips" gave them no additional information, (2) participants' default strategy, although different from the instructed one, led to similar performance levels, or, (3) the default strategy—while not optimal—was so natural and intuitive, that participants were reluctant or unable to shift to a potentially better strategy. Related to these alternatives are the questions: What is the default strategy that people use with PECs? Why is this strategy natural? Why is this strategy not optimal, in the sense that it leads to less-than-perfect performance (e.g. compared to the NEC group of Exp. 3)? We examine two alternative strategies in light of these questions.

Similarity based strategy

PECs seem to be naturally suited to an exemplar-like strategy, based on the storage of number of examples, or to a prototype-like strategy, based on abstraction of typical class elements. In our setup, however, participants were shown only pairs of objects of the same class (PECs). It may be difficult to build a prototype from two examples, and may be even more difficult to use an exemplar-based strategy with only two exemplars per category. Furthermore, the chief (an exemplar from the target category) was not necessarily from one of the categories shown in the constraints—and in fact usually was not—so that participants had to decide who belongs to the chief's class based on only one example from the target category. Nevertheless, participants could derive the size and shape of typical classes (in the multi-dimensional space) by averaging over the pairs shown, and, using the chief as the prototype of this unknown class, decide which other objects belong to it. Thus, we cannot rule out the possibility that people use this type of strategy, which may be natural for PECs, even though it is not optimal in the current setting, which forces generalization using a rule-based strategy.

Rule based strategy

A strategy that is based on a rule determined by the constraints provided to the participant, could guarantee perfect performance in our experimental setting, if it were used correctly. Specifically, participants could reliably derive from PECs the identity of the dimensions that are relevant or irrelevant to classification in each experimental trial. They could do this in one of two ways: (a) For each pair in a PEC, find the dimension or dimensions that differentiate the two stimuli, and identify them as irrelevant. After all the irrelevant dimensions are collected (a union operation), identify the remaining dimensions as the relevant dimensions (a set-complement operation). This strategy is the one provided to participants in the directed highPEC condition of Experiment 3. (b) For each pair in a PEC, find the set of dimensions shared by the two examples, and identify these dimensions as potentially relevant. As additional constraint pairs are examined, compute the intersection of the identified sets of dimensions, i.e., the dimensions that are shared in all the pairs. The result is the set of relevant dimensions. If participants used one of these methods for the PEC condition, they could have ended up with an elevated level of False-Alarms (as seen even in "Experiment 3"), because they may have missed less-salient relevant dimensions either when performing the set complement operation (in method a) or initially in identifying all the similarities within a constrained pair (method b). This error, due to missing less-salient dimensions, is prevalent in real-world cases, where the full group of possible dimensions may not be known or even inferable. Thus, even using this optimal strategy for PECs does not guarantee perfect performance.

We return to the question raised at the beginning of this section: Why is there no improvement of performance in the PEC condition when directions are provided? We are left with the three possibilities outlined there, which we now express in terms of the two strategies outlined above: (1) Participants actually use the rule-based strategy from the outset, but this strategy does not lead to perfect performance. (2) They may intuitively use a similarity-based strategy and then indeed shift their strategy, but performance may not improve, since the False-Alarm level remains high. (3) Participants may intuitively use a similarity-based strategy, and, since this strategy is quite effective even when performing a rule-based task, they may be reluctant to learn another strategy, and thus do not shift to the rule-based strategy even when given directions for its use. This latter possibility is supported by earlier studies showing a tendency of participants to use similarity-based categorization strategies even when an explicit rule is provided (Allen and Brooks 1991).

Strategies for using negative equivalence constraints

We compare the cases of PECs and NECs in terms of the two strategies presented above. The use of an exemplar- or prototype-like strategy is even less appropriate to NECs than to PECs, and may be impossible even for highly informative NECs since this strategy is based on similarities among objects of the *same* class. An attempt to use this strategy with NECs must lead to very poor performance, similar to baseline performance with poorly-informative constraints or even no constraints at all. This is just what we found for many of our participants.

Alternatively, participants could use a rule-based strategy, parallel to the one suggested above for PECs. Participants would identify as relevant the single dimension differentiating the stimuli of each pair, and collect these (a union operation) to form the set of relevant dimensions. No additional set-complement operation is needed, and less-salient relevant dimensions are highlighted directly by the constraints provided. Thus, perfect performance—without elevated False-Alarms—is likely, once this strategy is known and used. This is what we found for some participants even without giving them directions (Exp. 2), and for all participants who were given directions (Exp. 3).

PECs versus NECs

Two additional differences between the use of PECs and NECs require clarification: (1) the individual differences in the use of NECs, leading to a non-uniform distribution in the use of highNECs versus the uniformity and unimodal distribution for highPECs, and (2) the usefulness of giving directions for use of NECs but not of PECs.

(1) The individual differences may be explained by two characteristics of the information provided by NECs, one which facilitates their use, and one which complicates their use: NECs provide information indicating a dimension that is *relevant* to categorization. Such information may be more easily integrated than that provided by PECs—which decisively indicate dimensions that are irrelevant. On the other hand, NECs provide information regarding two categories, both of which must be kept in mind simultaneously. This may be more difficult than the use of PECs, which relate to one category at a time. Thus, use of NECs inherently contains both a difficult aspect (relating to two categories simultaneously) and an easy aspect (directly pinpointing relevant dimensions). The relative weight of these two factors may depend on the strategy each participant implements (e.g. rule vs. similarity based), leading to the non-unimodal distribution in their use, and explaining why when not provided with additional directions, only some participants effectively used highly informative NECs.

(2) The second major difference between PECs and NECs is that performance with NECs, in contrast to PECs, benefited significantly from the instructions regarding the optimal categorization strategy. Two factors may have contributed to this difference: (1) Participants were more open to advice on how to use NECs because they did not have a strong intuitive idea of what to do a priori; they may have been using an ineffective exemplar- or prototype-like strategy, another strategy, or no strategy at all. (2) It was easier for the participants to learn the rule-based strategy with NECs, perhaps because it did not involve a set-complement operation. This latter ease of abstracting the optimal strategy may also be the source of the excellent

performance by the good-performing highNEC subgroup in Experiment 2; they may have used this strategy even without directions.

Our findings reveal one way that behavior reflects statistical properties of objects and categories in the world: People have the tools needed to integrate PECs in category learning, since PECs are generally informative. In the case of NECs, only some people have the wherewithal for proper use of this information. Others, who lack this ability, are able to acquire it when provided with the necessary directions. Interestingly, NECs lead to better performance, once the proper strategy is found naturally or through directions.

Implications for the categorization hierarchy

The current findings have important implications for understanding known phenomena in category learning, and may provide an effective tool for predicting performance in a variety of category-learning tasks. As an example, the tendency of children to over-generalize when classifying objects (Clark 1973; Neisser 1987) may be seen as a consequence of their using mostly PECs, which, as pointed out above, can lead to disregarding less-salient, but relevant dimensions and a subsequently higher rate of False-Alarms. Perhaps later in life over-generalization is reduced when more refined strategies are acquired and better dimension weighting is attained including less salient dimensions see (R. Hammer et al., submitted; Diesendruck et al. 2003; Hammer and Diesendruck 2005; Sloutsky 2003 for similar thoughts). For example, as we saw above, some people do learn to use NECs in the rare cases when they are informative, resulting in the reduction of such False-Alarms.

These findings also shed light on the hierarchical structure of our conceptual knowledge, pinpointing differences between levels and the possible source of the order of acquiring them. Specifically, superordinate and basic level categories are expected to contain objects which are both similar in many aspects (dimensions), but also are dissimilar in many other aspects (Neisser 1987; Murphy 2004; Rosch and Mervis 1975; note that superordinate categories require a further level of abstraction and use of more “functional” rather than perceptual dimensions than basic-level categories. As we demonstrated, identifying the relevant dimensions in such categorization scenarios can be done only from PECs—but not from NECs because two negatively constrained objects, i.e., from different categories, are expected to be dissimilar in many dimensions, only some of which are relevant. Informative NECs (with only very few discriminating dimension) should therefore be extremely rare for superordinate and basic level categories. The use of NECs might be relevant only on those occasions when a supervisor intentionally selects informative NECs or highlights relevant discriminating dimensions. For instance, an adult telling a child, “You see these two (pointing to a horse and a dog), they are not the same because this one is large and that one is small” is adding to the information of the constraint itself, highlighting size as a relevant dimension for discriminating dogs from horses and shifting the child’s attention from other irrelevant dimensions in which the regarded two instances differ. Similarly, training medical diagnosis is more effective when novices are provided with an explicit rule including a list of differentiating symptoms. Nevertheless, after encountering a sufficient number of exemplars, further improvement involves use of similarity-based strategies (Kulatunga-Moruzi et al. 2001).

The case of subordinate level categories is different. Here a pair of negatively constrained objects from different subordinate categories—but the same basic level category—will generally differ on very few dimensions that will also be less distinct (Murphy 2004). Such a constraint may often be informative. Thus, subordinate level categories can be learned from either PECs or NECs. Moreover, subordinate level PECs may be less useful since objects from subordinate categories are usually already perceived as similar, and so are likely to be perceived as belonging to the same category—as they are at the basic categorization level. In this context, PECs will not be useful in highlighting non-salient relevant dimensions for categorization although they might still be useful for identifying salient non-relevant ones.

On the other hand, NECs may help in breaking default beliefs about the relation between highly similar exemplars, as illustrated in Fig. 11: This example suggests that NECs may act as a useful tool in boosting perceptual learning or dimension-creation by directing attention to subtle differences, between constrained instances, that otherwise would be disregarded or overshadowed by more salient ones. Later, the importance for categorization of these newly learned dimensions can be further evaluated. Similar ideas for NECs playing such a function are implied by Schyns et al. (1998) who discussed diagnostic-driven learning and differentiation in supervised categorization tasks. They and others provided examples for sensitization effects occurring only on task relevant dimensions that were identified via training in supervised categorization (Goldstone 1994b) or similarity judgment (Livingston et al. 1998) tasks.

These differences in the possible roles of PECs and NECs for learning different levels of the categorization hierarchy may explain why it is often hard to learn subordinate level categories. PECs suffice—and may even be better—for learning basic level and superordinate level categories, but NECs may be crucial for learning

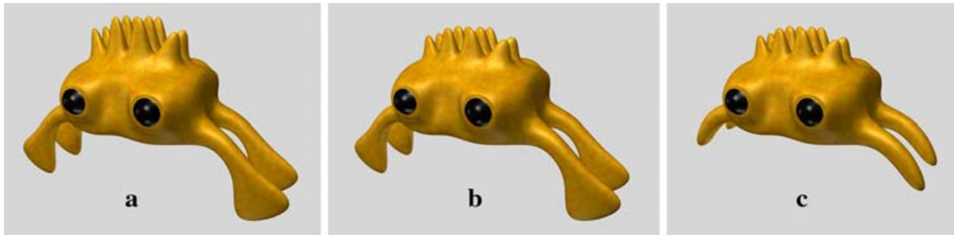

Fig. 11 The role of NECs in perceptual learning and category learning: *Before reading any further*—try to determine quickly which of the three creatures above belongs to a different species than the other two. When asked, most people first choose creature C as the different one since its limbs are very different than those of the other two. But when provided with an indication that creature A is not of the same kind as creature B (NEC), people become aware of the differences between the creatures in terms of the size of the spikes on their back, a dimension that was previously left unnoticed. When provided instead with the corresponding PEC, indicating that B and C are from the same category, people understand that the limb shape is not important, but yet they still do not notice or identify the spikes-size dimension as a relevant one. We claim that using NECs in such a context is essential for learning. Our current findings suggest that even when overcoming the perceptual limitations when provided by such NECs, many will still find it difficult to correctly use the information provided by them

subordinate categories. Therefore, our current findings, demonstrating difficulties in using NECs even in a simple categorization task using easily identified dimensions (as verified in Exp. 3), suggest that subordinate-level categorization will be hard even without perceptual difficulties. Expertise must involve not only better perceptual capabilities, but also an ability to implement an appropriate strategy for using NECs. Viewed from a different perspective, perhaps the fact that discriminating subordinate level categories is less frequently necessary in everyday life may underlie people's difficulty in using NECs.

Other theoretical implications

The current findings may have implications for additional category learning phenomena. More specifically, the role of PECs versus NECs may change when faced with complex or fuzzy boundaries including boundaries in typical XOR learning e.g. (Dixon et al. 2000; Kinder and Lachnit 2003; Palmeri and Noelle 2002) or information-integration tasks e.g. (Ashby and Maddox 2005; Avrahami et al. 1997; Palmeri and Noelle 2002). The conclusions discussed earlier are consistent with NECs being more suitable for these difficult cases, since they may more clearly define questionable boundaries. We suggest that these cases may be more difficult also because they depend on the use of NECs.

Although the current research was designed to test human use of equivalence constraints in category learning, it also raises theoretical issues that are directly relevant to other fields of research. In the introduction, we described theoretical limitations in the use of PECs and NECs that are relevant in any context involving the identification of common or discriminating attributes in a multidimensional object space. The differentiating properties of PECs and NECs should affect their use by any agent, whether human, animal or machine, when faced with a category-learning, discrimination, or similarity-judgment task.

For example, in many studies involving animal training, the often wearying effort of teaching an animal to discriminate between multidimensional stimuli e.g. (Brosch et al. 2005) can be avoided by the use of well chosen stimulus pairs during training. The use of highly informative NECs—when possible—is expected to be most beneficial for teaching discrimination between different types of stimuli. At the same time, it would be of interest to determine if animals possess biases against the use of even informative NECs, similar to those observed here in humans.

Similarly, in the context of machine learning, it has already been demonstrated that an EM (Expectation-Maximization) clustering algorithm designed for using equivalence constraints has difficulty using even informative NECs, but easily succeeds in learning target categories when provided with PECs (Hammer et al. 2007; Hertz et al. 2003; Shental et al. 2004). This limitation arises from the fact that this algorithm represents categories by cluster centers and the distributions around these centers, i.e., they are conceptually similar to prototype-based classifiers. As described above, PECs are more efficient than NECs for calculating prototypes; but see (Winston 1982 on learning from "near misses", as an example of a possible algorithm which learns from NECs).

Future research

The current study provides insight into category learning strategies and dynamics. Further study is needed to address related questions concerning the separate role of PECs and NECs. As discussed earlier, we expect differences in the use of PECs and NECs in early development: Children might use PECs as do adults, but their use of highly informative NECs may be similar to the poorly performing adult group (Experiment 2). Recent findings already provide some support for this claim (Hammer et al., submitted). In another direction, it would be interesting to see what happens when non-binary dimensions discriminate between otherwise very similar categories (i.e., with similar property values). Here informative NECs may play a more significant role, as they do for subordinate categories (see above).

Acknowledgments This study was supported by a "Center of Excellence" grant from the Israel Science Foundation, a grant from the US-Israel Binational Science Foundation, and a grant by the EU under the DIRAC integrated project IST-027787. Preliminary results of this study were presented in the annual meeting of the Cognitive Science Society, Stresa, Italy, July 2005. We would like to thank Lee Brooks and Gil Diesendruck for their comments. We also thank Michael Ziessler and an anonymous reviewer for their useful comments.

Appendix 1

We analyze the dependence of the number of possible PECs, NECs, and highNECs on the number of objects and categories. Note that all PECs are informative for identifying relevant dimensions while in the case of NECs, only the highNECs (negative constraints made up of two objects from two different categories that differ in their value on only a single dimension) are adequately informative for such a task. To simplify the discussion, we assume that the number of objects in each category is identical.

Specifically, let

$c =$ the number of categories.
$n =$ the number of objects in each category.
$d =$ the number of relevant dimensions, assuming binary dimension, $d = \log_2 c$

It follows that

$$\text{number of PEC} = \frac{nc(n-1)}{2};$$

$$\text{number of NEC} = \frac{n^2c(c-1)}{2};$$

$$\text{number of highNEC} = \frac{\text{ncd}}{2}.$$

This calculation shows that the total number of PECs is much smaller than the total number of NECs specifically when the number of categories, c, is large. In addition, highNECs (NECs which provide 1 Bit of information) are a small subset of NECs when the number of category members, n, is large. Specifically:

$$\frac{\text{PEC}}{\text{NEC}} = \frac{nc(n-1)}{n^2c(c-1)} \approx \frac{(n-1)}{n(c-1)} \approx \frac{1}{c} \ll 1$$

$$\frac{\text{highNEC}}{\text{NEC}} = \frac{\text{ncd}}{n^2c(c-1)} \approx \frac{\log_2 c}{nc} \ll \frac{1}{c}$$

In the current experiment, $nc = 32$. When $d = 2$, $c = 4$ and $n = 8$. Then, there are 112 PECs and 384 NECs, of which 32 are highNEC. When $d = 3$, $c = 8$ and $n = 4$. Then, there are 48 PECs and 448 NECs, of which only 48 are highNEC.

References

Allen SW, Brooks LR (1991) Specializing the operation of an explicit rule. J Exp Psychol Gen 120:3–19

Ashby FG, Maddox WT (2005) Human category learning. Annu Rev Psychol 56:149–178

Avrahami J, Kareev Y, Bogot Y, Caspi R, Dunaevsky S, Lerner S (1997) Teaching by examples: Implications for the process of category acquisition. Q J Exp Psychol 50A(3):586–606

Brooks LR, Norman GR, Allen SW (1991) The role of specific similarity in a medical diagnostic task. J Exp Psychol Gen 120:278–287

Brosch M, Selezneva E, Scheich H (2005) Nonauditory events of a behavioral procedure activate auditory cortex of highly trained monkeys. J Neurosci 25(29):6797–6806

Clark HH (1973) Space, time, semantics, and the child. In: Moore TE (ed) Cognitive development and the acquisition of language. Academic Press, New York

Cohen AL, Nosofsky RM (2000) An exemplar-retrieval model of speeded same-different judgments. J Exp Psychol Hum Percept Perform 26:1549–1569

Diesendruck G, Hammer R, Catz O (2003) Mapping the similarity space of children and adults' artifact categories. Cogn Dev 118:217–231

Dixon MJ, Koehler D, Schweizer TA, Guylee MJ (2000) Superior single dimension relative to "exclusive or" categorization performance by a patient with category-specific visual agnosia: empirical data and an ALCOVE simulation. Brain Cogn 43(1–3):152–158

Garner W (1978) Aspects of a stimulus: features, dimensions and configurations. In: Rosch E, Lloyd B (eds) Cognition and categorization. Lawrence Erlbaum, Hillsdale

Gentner D, Kurtz K (2005) Learning and using relational categories. In: Ahn WK, Goldstone RL, Love BC, Markman AB, Wolff PW (eds) Categorization inside and outside the laboratory. APA, Washington, DC

Goldstone RL (1994a) The role of similarity in categorization: providing a groundwork. Cognition 52:125–157

Goldstone RL (1994b) Influences of categorization on perceptual discrimination. J Exp Psychol Gen 123(2):178–200

Grier JB (1971) Nonparametric indexes for sensetivity and bias: Computing formulas. Psychol Bull 75:424–429

Hammer R, Diesendruck G (2005) The role of dimensional distinctiveness in children's and adults' artifact categorization. Psychol Sci 16(2):137–144

Hammer R, Hertz T, Hochstein S, Weinshall D (2007) Classification with positive and negative equivalence constraints: theory, computation and human experiments. In: Mele F, Ramella G, Santillo S, Ventriglia F (eds) Brain, vision, and artificial

intelligence: second international symposium, BVAI 2007. Lecture notes in computer science. Springer, Heidelberg, pp 264–276
Hammer R, Diesendruck G, Weinshall D, Hochstein S. The development of category learning strategies: what makes the difference? (submitted)
Hertz T, Shental N, Bar-Hillel A, Weinshall D (2003) Enhancing image and video retrieval: learning via equivalence constraints. IEEE Conference on computer vision and pattern recognition, Madison WI, June 2003
Huettel SA, Lockhead GR (1999) Range effects of an irrelevant dimension on classification. Percept Psychophys 61(8):1624–1645
Jones M, Love BC (2004) Beyond common features: the role of roles in determining similarity. Proceedings of the cognitive science society
Jones M, Love BC, Maddox WT (2006) Recency as a window to generalization: separating decisional and perceptual sequential effects in category learning. J Exp Psychol Learn Mem Cogn 32:316–332
Kareev Y, Avrahami J (1995) Teaching by examples: the case of number series. Br J Psychol 86:41–54
Katz JJ, Postal PM (1964) An integrated theory of linguistic descriptions. MIT Press, Cambridge
Kinder A, Lachnit H (2003) Similarity and discrimination in human Pavlovian conditioning. Psychophysiology 40:226–234
Klayman J, Ha Y-W (1987) Confirmation, disconfirmation and information in hypothesis testing. Psychol Rev 94:211–228
Kulatunga-Moruzi C, Brooks LR, Norman GR (2001) Coordination of analytic and similarity-based processing strategies and expertise in dermatological diagnosis. Teach Learn Med 13(2):110–116
Levine M (1966) Hypothesis behavior by humans during discrimination learning. J Exp Psychol 71:331–338
Livingston KR, Andrews JK, Harnad S (1998) Categorical perception effects induced by category learning. J Exp Psychol Learn Mem Cogn 24:732–753
Medin DL, Schaffer MM (1978) Context theory of classification learning. Psychol Rev 85:207–238
Medin DL, Goldstone RL, Gentner D (1993) Respect for similarity. Psychol Rev 100(2):254–278
Mooney RJ (1993) Integrating theory and data in category learning. In: Nakamura GV, Taraban R, Medin DL (eds) The psychology of learning and motivation: categorization by humans and machines, vol 29. Academic Press, San Diego, pp 189–218
Murphy G (2004) The big book of concepts. MIT Press, Cambridge
Murphy G, Medin DL (1985) The role of theories in conceptual coherence. Psychol Rev 92:289–316
Neisser U (1987) Concepts and conceptual development. Cambridge University Press, Cambridge
Nosofsky RM (1987) Attention and learning processes in the identification and categorization of integral stimuli. J Exp Psychol Learn Mem Cogn 13:87–108
Nosofsky RM (1988) Similarity, frequency, and category representations. J Exp Psychol Learn Mem Cogn 14:54–65
Nosofsky RM (1990) Relation between exemplar-similarity and likelihood models of classification. J Math Psychol 34:812–835
Ohl FW, Scheich H, Freeman WJ (2001) Change in pattern of ongoing cortical activity with auditory category learning. Nature 412:733–736
Palmeri TJ, Noelle D (2002) Concept learning. In: Arbib MA (ed) The handbook of brain theory and neural networks. MIT Press, Cambridge
Rosch E, Mervis CD (1975) Family resemblance studies in the internal structure of categories. Cogn Psychol 7:573–605
Rouder JN, Ratcliff R (2006) Comparing exemplar- and rule-based theories of categorization. Curr Dir Psychol Sci 15:9–13
Schyns PG, Goldstone RL, Thibaut JP (1998) The development of features in object concepts. Behav Brain Sci 21:1–54
Shental N, Bar-Hillel A, Hertz T, Weinshall D (2004) Computing Gaussian mixture models with EM using equivalence constraints. In: Proceedings of neural information processing systems, NIPS 2004
Shepard RN, Hovland CL, Jenkins HM (1961) Learning and memorization of classifications. Psychol Monogr 75:1–42
Sloutsky VM (2003) The role of similarity in the development of categorization. Trends Cogn Sci 7:246–251
Smith EE, Medin DM (1981) Categories and concepts. Harvard University Press, Cambridge
Stanislaw H, Todorov N (1999) Calculating signal detection theory measures. Behav Res Methods Instrum Comput 31(1):137–149
Stewart N, Brown GDA (2005) Similarity and dissimilarity as evidence in perceptual categorization. J Math Psychol 49:403–409
Stewart N, Brown GDA, Chater N (2005) Absolute identification by relative judgment. Psychol Rev 112:881–911
Tversky A (1977) Features of similarity. Psychol Rev 84(4):327–352
Tversky A, Gati I (1982) Similarity, separability, and the triangle inequality. Psychol Rev 89:123–154
Wason PC (1960) On the failure to eliminate hypotheses in a conceptual task. Q J Exp Psychol 12:129–140
Whitman JR, Garner WR (1962) Free recall learning of visual figures as function of form of internal structure. J Exp Psychol 64(6):558–564
Winston PH (1982) Learning by augmenting rules and accumulating censors, Memo 678, MIT AI Lab, May 1982