Acta Psychologica 130 (2009) 81–94

Contents lists available at ScienceDirect

Acta Psychologica

journal homepage: www.elsevier.com/locate/actpsy

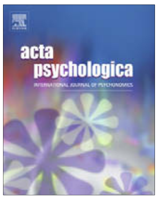

Situational information contributes to object categorization and inference

Sergio E. Chaigneau [a,*], Lawrence W. Barsalou [b], Mojdeh Zamani [c]

[a] *Escuela de Psicología, Universidad Adolfo Ibáñez, Avenida Diagonal Las Torres 2640, Peñalolén, Santiago, Chile*
[b] *Department of Psychology, Emory University, Atlanta, Georgia 30322, United States*
[c] *France Télécom R&D, DIH/UCE, France*

ARTICLE INFO

Article history:
Received 4 June 2008
Received in revised form 9 October 2008
Accepted 13 October 2008
Available online 28 November 2008

PsycINFO classification:
2340

Keywords:
Situated cognition
Categorization
Events
Artifacts

ABSTRACT

Three experiments demonstrated that situational information contributes to the categorization of functional object categories, as well as to inferences about these categories. When an object was presented in the context of setting and event information, categorization was more accurate than when the object was presented in isolation. Inferences about the object similarly became more accurate as the amount of situational information present during categorization increased. The benefits of situational information were higher when both setting and event information were available than when only setting information was available. These findings indicate that situational information about settings and events is stored with functional object categories in memory. Categorization and inference become increasingly accurate as the information available during categorization matches situational information stored with the category.

1. Introduction

Categorization research focuses on the physical properties of objects. In studies of category learning, participants typically learn how the physical properties of objects—shape, color, material, texture, etc.—predict the categories being learned. One central issue in the categorization literature concerns whether people use rules, prototypes, or exemplars to represent the object properties that predict categories. Other important issues concern the attentional weighting of object properties, the differential roles of analytic and holistic properties, the effects of correlated properties, and the distribution of object properties across taxonomic levels. Clearly, the extensive research that has investigated how people represent and process object properties has taken the understanding of categorization forward considerably. Nevertheless, research on object categorization has ignored other potentially interesting factors in categorization, in particular, the role of situational information.

Do the situations in which objects typically occur play a role in categorization? Are these situations part of these objects' conceptual content? We propose that object concepts not only contain situational content, but that this content plays central roles in object categorization. We define a situation as a setting that contains a focal object often involved in an event over a temporal duration (Yeh & Barsalou, 2006). We further define a setting as a place where a focal object is generally found, the set of associated objects that typically populate the place, and the relations that usually hold between these entities (e.g., spatial relations). Finally, we define an event as a dynamic process extended over time in which the focal object and associated objects participate.

For decades, researchers across multiple disciplines have argued that situations play central roles throughout cognition. During language comprehension, a text can be incomprehensible if the relevant situation is not apparent (e.g., Bransford & Johnson, 1973). During conversation, shared situations help human speakers establish common ground (e.g., Clark, 1992); shared situations also support non-human communicators (e.g., Smith, 1977). Extensive evidence demonstrates that situation models underlie people's representations of text meaning (e.g., Sanford & Garrod, 1981; Zwaan & Radvansky, 1998). In general, language comprehension appears to be a heavily situated process (Barsalou, 1999a). During problem solving and reasoning, drawing valid conclusions without the support of concrete situations is often difficult (e.g., Cheng & Holyoak, 1985; Gick & Holyoak, 1980; Johnson-Laird, 1983). In cognitive development, situations are central to acquiring cognitive and social skills (e.g., Vygotsky, 1991). During social interaction, situations are central to predicting behavior (e.g., Mischel, 1968; Smith & Semin, 2004). In linguistics, situations underlie the theory of construction grammar, with the content and relations of syntactic structures evolving out of the analogous structure in situations (e.g., Goldberg, 1995). In philosophy, situations motivated the theory of situation semantics, with logical inference being optimized

---

[*] Corresponding author. Tel.: +56 2 675 4436.
E-mail address: sergio.chaigneau@uai.cl (S.E. Chaigneau).

0001-6918/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.actpsy.2008.10.004

in the context of specific situations (e.g., Barwise & Perry, 1983). In artificial intelligence, situating robotic cognition in physical environments greatly enhances practical intelligence (Brooks, 1991; Kirsh, 1991). At a broader level, general arguments about the centrality of situations in cognition have been presented by Clark (1997), Dunbar (1991), Glenberg (1997), and Greeno (1998). The Cambridge Handbook of Situated Cognition further documents the broad roles of situations throughout cognition (Robbins and Aydede, 2009).

Given the breadth of situational effects, it would be surprising not to find them in object categorization. Yeh and Barsalou (2006) argue that situations are central to the representation and processing of categories (also see Barsalou, 2003b; 2005b; in press; Barsalou, Niedenthal, Barbey, & Ruppert, 2003). They review extensive evidence showing that categories are stored with situational knowledge, and that this knowledge plays important roles in a variety of conceptual activities. Most of this evidence, however, comes from tasks other than categorization (e.g., comprehension, memory), with only modest evidence coming from research on object recognition (e.g., Bar & Ullman, 1996; Biederman, 1972; 1981; Dobel, Gumnior, Bölte, & Zwitserlood, 2007; Palmer, 1975). These latter studies do indeed demonstrate that settings affect categorization. When an object is perceived in a typical setting, it is categorized more easily than when it is not (e.g., categorizing a stove in a kitchen vs. a park). Interestingly, findings like these have been largely ignored in subsequent research on object categorization. Our aim here is to fill this gap by providing evidence that situations are central to categorization.

As will become clear next, functional objects offer an ideal test bed for demonstrating situational effects on object categorization. Representing the function of an object requires extensive use of situational information. Although physical properties are important, many other kinds of properties are important as well. According to Barsalou, Sloman, and Chaigneau (2004) HIPE theory of function, an adequate representation of an object's function must contain information about History, Intention, Physical structure, and Events (also see Wimsatt, 1972; Wright, 1973). Knowing the historically intended function of an object is important for understanding its function. Knowing the intention (goal) that an agent is trying to accomplish with the object is also important. Knowing the object's physical parts, as well as its physical setting, is essential. Knowing an agent's actions and how the object behaves as a result of these actions is also essential.

To make these four kinds of information in the HIPE theory more concrete, consider how they apply to the concept of hammer. To fully understand the function of a hammer, knowing that it was historically intended for pounding, is important, as is knowing that an agent currently intends to use it for this purpose. Knowing the physical structure of the hammer (e.g., metal head, wooden handle) and the physical structure of related setting objects (e.g., nails, boards) is also important. Finally, knowing that an agent swings a hammer into contact with another object is important, and also understanding that when the other object is hit, it tends to move or shatter.

As this example illustrates, all this information is important for representing function. Function cannot be reduced to a subset. In particular, the function of an object cannot be reduced to its physical properties alone. Instead, people only fully understand its function once they represent all this information, which includes the situations in which the object is perceived. Chaigneau, Barsalou, and Sloman (2004) provided initial evidence for the HIPE theory.

Developmental research on the categorization of functional objects similarly demonstrates that situational information is crucial when categorizing these objects. If a perceived object affords a function associated with a category, perceiving the function provides evidence that the object belongs to the category. Infants as young as 14 months use perceived function—not only physical properties—in categorization (Madole & Oakes, 2004). Children as young as 2 and 3 commonly use function when categorizing physical objects (Kemler-Nelson, 1999; Kemler-Nelson, Russell, Duke, & Jones, 2000). Functional reasoning becomes increasingly important for categorization across development in older children and adults (e.g., Kemler-Nelson, 1995; Tomikawa & Dodd, 1980).

1.1. The basic paradigm

The three experiments reported here assessed the role of situational information in adult categorization and inference. To address this issue, we developed a paradigm that had the following features. First, the objects that participants categorized belonged to familiar object categories, thereby allowing us to assess how natural knowledge affects categorization. Second, the objects were novel, such that participants had no previous experience with these particular instances of the categories. Novel objects were necessary to avoid ceiling effects. Had we used everyday objects, participants would have easily categorized them based on superficial diagnostic features, and situations would have become uninformative. Third, the categories used were functional objects, thereby ensuring that situational information would be potentially relevant for their categorization. The functional object categories used were similar to those used previously in research with children (e.g., Kemler-Nelson, 1995; 1999; Kemler-Nelson et al., 2000) and with non-human primates (e.g., Hauser, 1997; Hauser, Pearson, & Seelig, 2002). Fourth, as in the previous research with children and non-human primates, we used actual physical objects, rather than property lists or computer images. Using actual objects allowed us to assess the role of situational information in the kinds of physical settings, where people typically categorize objects. Furthermore, psycholinguistics research has demonstrated that using actual physical settings produces powerful situational effects similar to those of interest here (e.g., Tanenhaus, Spivey-Knowlton, Eberhard, & Sedivy, 1995).

To get a sense of the paradigm, perform the following exercise. First, without looking, cover up both Panel B at the bottom of Fig. 1, and also the figure caption below it. Then, look at the 15 objects above in Panel A. See if you can determine the category of each object. Also, what objects do you think belong to the same category? Finally, while keeping the figure caption covered at the bottom, uncover Panel B above, and look at the three assembled systems. What category do you think each system belongs to?

Although not obvious, each system in Panel B is a catapult that flings a projectile toward a target. As Panel A illustrates, each system contains five novel objects, shown assembled in Panel B. Furthermore, the three objects in each row of Panel A share belong to the same (familiar) category for a functional object. Specifically, each row shows three objects that belong to the functional categories of fulcrum, lever, weight, projectile, and target, respectively. The fulcrum supports its lever and allows the lever to pivot. The lever holds both the weight and projectile in place simultaneously. When the weight is placed on its lever, it provides the system with the energy to fling the projectile. When the agent releases the lever, it swings downward under the force of the weight, flinging the projectile toward the target, which serves as the reference point through which the projectile, ideally, should travel.

Three levels of situational information were manipulated in the experiments to follow: object information alone (the object condition), object plus setting information together (the setting condition), and object plus setting and event information (the event condition). Each condition and its predictions are discussed in turn.

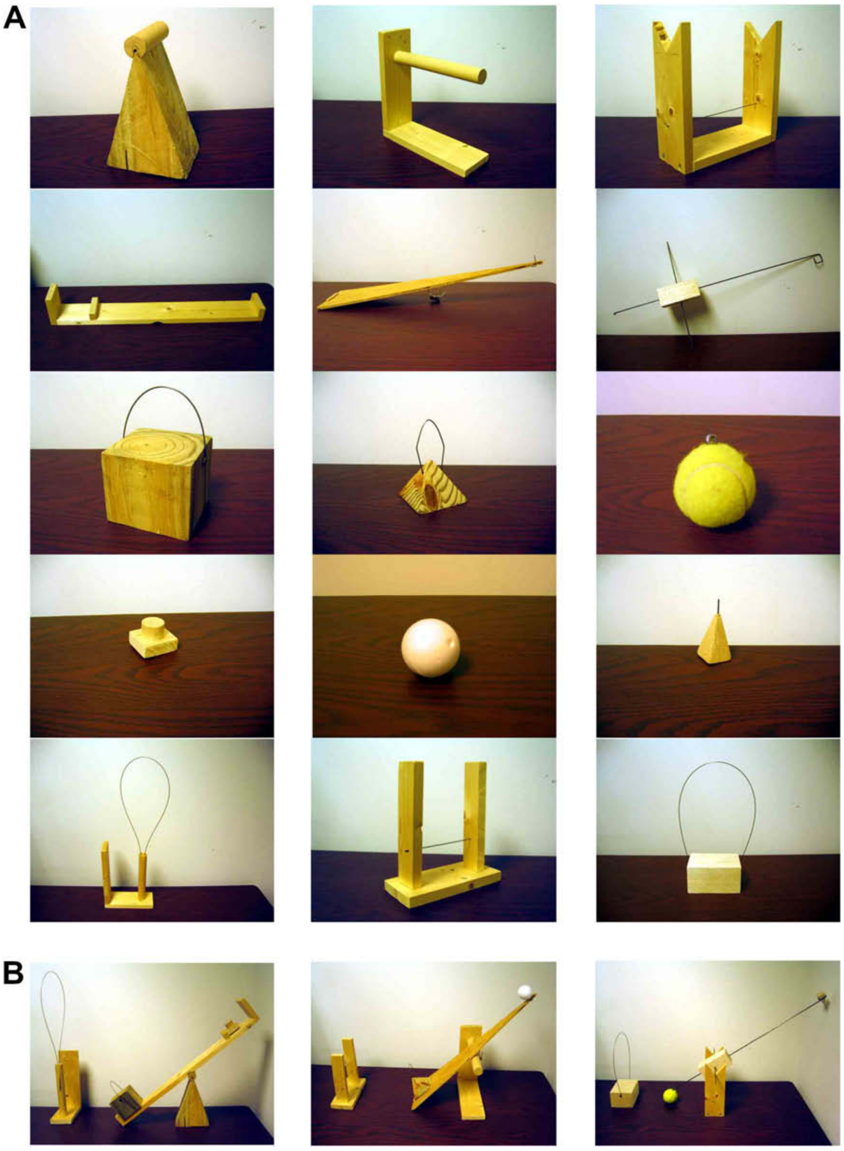

Fig. 1. Panel A. The five functional objects in each of the three systems. The objects from Systems 1, 2, and 3 are shown in columns from left to right, respectively. The fulcrums, levers, weights, projectiles, and targets are shown in the rows from top to bottom, respectively. Panel B. Each system assembled, as presented in the setting condition.

Participants in the object condition only received information about each object's physical properties. Participants did not receive setting or event information. Specifically, participants viewed each actual object of a single system in Fig. 1A in isolation, not moving or performing its function. In most experiments, participants only viewed the five objects for a single system one at a time in this manner (i.e., they did not see objects from the other two systems).

If object properties are sufficient to categorize the objects from a system when studied individually, then participants should assign the a priori category to each object with a high level of accuracy. If the physical properties of an object afford its function (Gibson, 1979), then no setting or event information should be necessary for correct categorization.[1] Object properties should be sufficient. If, however, setting and event information are central to the a priori functional categories, then object properties alone should be insufficient for correct categorization. No category should become active, or several partially consistent categories should become active. Because situational information is important for categorization, participants should exhibit high levels of inaccuracy and uncertainty.

Participants in the setting condition received setting information about each object in the context of its assembled system. Specifically, participants viewed one of the assembled systems depicted in Fig. 1B, such that each object was categorized within the setting of its larger system. Notably, however, the system was static. No agent acted on it, nor did it move.

If situational information is central to the a priori categories, then setting information—the configuration of objects—should improve performance relative to the object condition. Because perceived setting information matches setting information stored in memory for the a priori categories, activating these categories should become more likely. Because no events are shown, however, the object and setting information perceived may still not be sufficient to activate the correct category for an object. To the extent that participants can simulate relevant events for the object,

---

[1] In principle, a Gibsonian affordance could be learned. However, in keeping with the strongest formulation of the concept, here we will use the term *affordance* to refer to possibilities for action that are perceived directly, without the mediation of representations.

however, this should increase the likelihood of correct categorization.

Participants in the event condition viewed an assembled system (as in the setting condition), and also watched an agent act on the system to produce its subsequent behavior. Thus, these participants saw each component of a system behave in a particular manner, together with the behaviors of the other components. Nothing was said about the agent's intention for using the system.

If situational information is central to the a priori categories, then event information—the behaviors of the objects—should improve performance relative to the setting and object conditions. Because perceived setting and event information matches setting and event information stored in memory for the a priori categories, activating these categories should be even more likely than in the setting condition. Because events are shown, perceived information for the object may optimally match information stored with the categories in memory. If all this situational information is central to the representations of these categories, then categorization may only occur at high levels of accuracy when all this information is present and matches these representations.

In summary, we assume that participants attempt to categorize the novel objects that they perceive in these experiments by instantiating familiar object categories. Because situational information represents these categories—not just object properties—matching these situational representations determines the success of categorization. As more situational information is presented, it increasingly matches representations of the relevant categories in memory, such that categorization becomes increasingly successful.

1.2. Overview of the experiments

Experiment 1 assessed whether object properties are sufficient to categorize the 15 objects in Panel A of Fig. 1, or whether situational information is necessary. If situational information is central for these categories and is necessary for categorizing their members, then presenting situational information should improve categorization performance. Furthermore, participants should categorize with high levels of accuracy when both setting and event information are present.

Experiments 2 and 3 shifted the focus from categorization to inference. Rather than clustering objects, participants were asked to generate properties for them. In Experiment 2, participants were asked to generate the function of each object (and also of the overall system). As more situational information becomes available during categorization, accessing the correct categories should become more likely, such that increasingly accurate functional inferences result.

In Experiment 3, participants were asked to "generate things that are typically true of objects of this type." These inference instructions were more open-ended than those for Experiment 2 that focused on function. Again, we predicted that the properties generated for a category should become increasingly accurate descriptions of its members as more situational information is available, thereby activating the correct a priori categories. Experiment 3 also assessed the impact of situational information on relational properties. Because relations are central to organizing knowledge about situations, the generation of relational properties should increase as situational information improves accuracy. Finally, Experiment 3 assessed whether inferential competence improves after seeing two examples of each category as opposed to one. If receiving setting and event information for one object is sufficient to access its category, then seeing the second example of the category should provide no added benefit.

At this point, one might worry that the three conditions described above (object, setting, event), only differ in the amount of information available to participants. From this perspective, it would be unsurprising to find that the more participants learned about objects in Fig. 1, the better they became at categorizing them. Note, however, that increasing accuracy was not our only prediction. Based on the situational content in the HIPE theory of function (Barsalou et al., 2004), we also predicted that specific types of properties would become more or less salient across our three conditions. As will become clear later, the distribution of property types produced by participants changed as a function of the three conditions, and did so in forms consistent with our theoretical analysis. Our three conditions were designed to activate different patterns of information, and it is these different patterns that explain increasing accuracy. Furthermore, one can imagine adding random types of information to object memories and not seeing cumulative benefits, because the additional information is irrelevant. Thus, a general increase in information cannot explain the predicted effects. Instead, as the HIPE theory predicts, increasing information only produces benefits when it resides in memories of the objects, and thus can facilitate retrieving them.

2. Experiment 1

This experiment assessed whether providing participants with setting and event information during categorization increases the activation of the a priori categories. To establish a measure of sorting accuracy with respect to the a priori categories, Experiment 1 used a cued sorting procedure. At the start of the experiment, five "cue" objects from one catapult system in Fig. 1 were laid out on a table in front of a participant (i.e., the fulcrum, a lever, a weight, a projectile, and a target from one particular system). The objects from another system were then placed separately on the table in a random arrangement. The participant's task was to sort each of the latter objects with the cue object that was most similar, thereby creating five clusters of two objects each. The participant then received the five objects from the third system in a random order, and had to sort each object into the most similar existing cluster of two objects.

If participants sorted perfectly according to the a priori categories, their five final clusters of three objects should have corresponded to the categories of fulcrums, levers, weights, projectiles, and targets. When participants sorted in this manner, they received a score of 10, indicating that they sorted all 10 objects correctly into the a priori categories. Participants were also asked to describe their sortings, which were coded for content.

2.1. Method

2.1.1. Participants and design

Participants were 36 Emory undergraduates, who participated for course credit (30 females, 6 males). The manipulation of situational information consisted of two levels: object vs. event (i.e., the setting condition was not included but was included in the next two experiments). Six participants were nested in each cell of the 2×3 design created by crossing situational information (object, event) and cuing system (1, 2, 3).

2.1.2. Materials and procedure

The materials were the three catapult systems shown in Fig. 1. Each system served as the cuing system for 12 of the 36 participants. The ten objects from the remaining two systems were laid out in two rows, one for each system, with the order of the rows, and the order of objects within rows, determined randomly for each participant. Prior to the participant entering the room, a sheet covered each of these two rows, so that participants could not see the objects underneath.

In the event condition, the cuing system was demonstrated twice, providing participants with setting and event information. The system was then left assembled, so that setting information was preserved for the duration of the experiment, to help participants remember how the system worked. Five pictures, one of each object in the cuing system, were then laid out in a random row in front of the system to serve as anchors for clusters. Pictures were used instead of the actual objects so that the system could remain assembled as participants sorted objects from the other two systems.

The first covered row of objects was uncovered, and participants were asked to place each of the five objects with the pictured object that best represented "the same type of object." The second covered row of objects was then uncovered, and participants placed each object with the cluster that best represented the same type. At the end of the procedure, five clusters existed, each containing a picture from the cuing system, and one object from each of the two sorted systems.

The procedure was identical in the object condition except that participants saw no demonstrations of the cuing system, such that they did not receive setting or event information. Instead, the five objects from the cuing system were laid out randomly in a row. Their pictures were also placed with them, so that the cues contained both actual objects and pictures, as in the event condition. Participants then sorted objects from the two covered rows, one system at a time, into three-object clusters of the same type.

Following sorting, participants in both groups provided a description for each of the five clusters.

2.2. Results

2.2.1. Cluster accuracy

The proportion of correct functional sorts out of 10 possible sorts was computed for each participant. A correct functional sort was defined as placing an object with the cued object having the same a priori function (i.e., fulcrum, lever, weight, projectile, or target). These proportions were transformed using the arcsin function and submitted to a 2 (situational information) × 3 (cuing system) between-participants ANOVA.

As Fig. 2A illustrates, participants in the event condition (62%) were much more accurate in clustering the objects than were participants in the object condition (29%) ($F(1, 30) = 10.21$, $MSe = 0.18$, $p < 0.01$, $R^2 = 0.25$, power = 0.87). There was no effect of cuing system ($F(2, 30) = 0.83$, $MSe = 0.18$, $p < 0.5$, $R^2 = 0.05$, power = 0.18), nor an interaction ($F(2, 30) = 2.39$, $MSe = 0.18$, $p < 0.25$, $R^2 = 0.14$, power = 0.45). Situational information had a substantial impact on clustering. When situational information was available, participants were twice as likely to access the a priori categories as when only physical properties were available.

2.3.1. Description accuracy

The first author rated each of the five descriptions that a participant provided (one per cluster) on a 1 to 3 scale (1 being completely inaccurate, 3 being perfectly accurate). To prevent bias, information about each participant's experimental condition was removed from his or her data file prior to coding. Perfect accuracy was defined as using the name of the correct a priori functional category for a cluster. Approximate descriptions received lower scores. To avoid problems that arose when objects having different a priori functions were sorted together, the correctness of the description was judged only with respect to the cue object—the other two objects were irrelevant.

To assess reliability, the descriptions of 30 participants (83% of the sample) were selected randomly and rated independently for description accuracy by a research assistant blind to the hypothesis and to participants' conditions. The Pearson correlation between ratings of the two judges was 0.97.

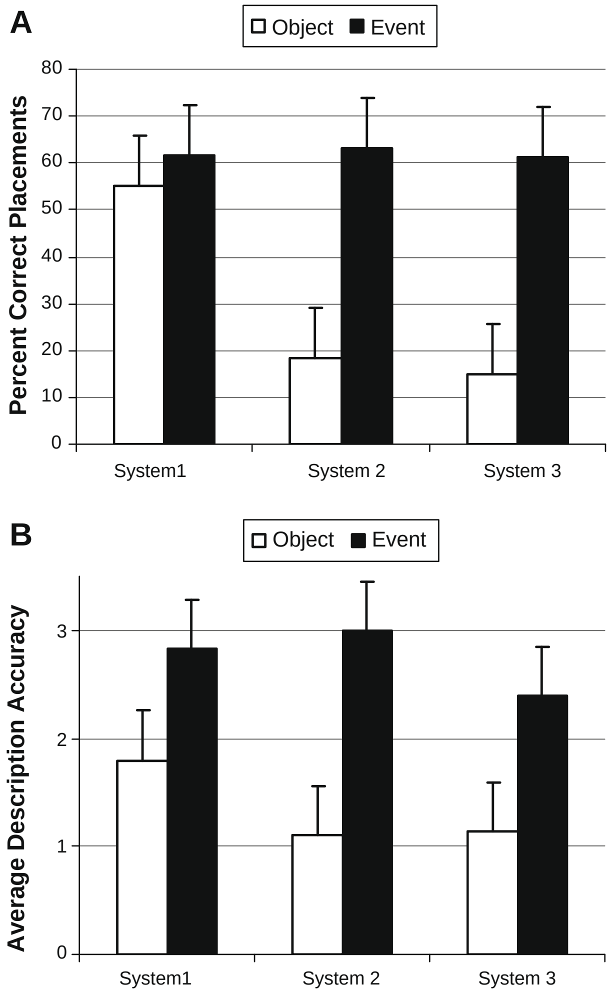

Fig. 2. Percent correct placement (Panel A) and average description accuracy (Panel B) in Experiment 1 as a function of situational information and cueing system. Error bars represent one standard error.

The average accuracy scores across clusters for each participant were submitted to a 2 (situation information) × 3 (cuing system) between-participants ANOVA. As Fig. 2B illustrates, the pattern of results for description accuracy was almost identical to the pattern for cluster accuracy. Participants in the event condition produced much more accurate descriptions (2.74) than did participants in the object condition (1.34), $F(1, 30) = 14.23$, $MSe = 1.24$, $p < 0.001$, $R^2 = 0.32$, power = 0.95. The overall analysis showed neither an effect of cuing system ($F < 1$), nor an interaction ($F < 1$).

The consistency of results from the accuracy measures for clusters and descriptions suggested they were closely related. Indeed, when the correlation between the two measures was assessed, it demonstrated a strong relationship between them ($r = 0.89$, $p < 0.001$).

2.4. Discussion

The presence of setting and event information dramatically altered sorting performance. When participants in the object condition lacked setting and event information, they did not organize clusters according to the a priori categories. Based on the results from a control experiment we performed, it is reasonable to infer that participants in Experiment 1 focused on physical object properties as they formed clusters. In that control experiment, when participants did not receive event information and were asked to freely cluster all 15 objects, their clusters and descriptions showed they overwhelmly relied on physical object properties to categorize. Conversely, when participants in Experiment 1 received setting and event information, they organized clusters around the a priori categories, and often described these clusters correctly. This finding indicates that setting and event information are stored in memory for the a priori categories, such that when participants received it prior to clustering, these categories became active and controlled categorization. Once these categories became available, they overrode the strong influence of object physical properties that determined the perceived similarity of objects in our control experiment.

3. Experiment 2

The first experiment addressed the role of situational information in object categorization. As we saw, the presence of situational information can have considerable impact on how people categorize objects. When situational information is present, it activates categories that remain inactive in its absence.

In the final two experiments, we added the setting condition. A possibility was that setting would provide sufficient information, and that learning about events is not necessary. The focus also changed from categorization to categorical inference. Rather than categorizing objects, participants generated properties for them. Generating properties constitutes a form of categorical inference because participants must produce properties that go beyond perceived object properties (Barsalou, 2003a; 2005a). In Experiment 2, participants attempted to generate the function of each object, and also the function of the overall system. In Experiment 3, participants were free to list a wider variety of properties for each object. Because generating accurate properties depends on accessing the correct category, properties should become increasingly accurate as more situational information is available.

Besides assessing the accuracy of functional inferences, Experiment 2 also assessed their content. Of interest was how the content of properties being generated changed as the amount of available situational information increased.

4.3. Method

4.3.1. Participants and design

Forty-five Emory undergraduates participated for course credit (31 females, 14 males). Two of the original participants were replaced because of experimenter error, and five because of equipment failure.

Five of the participants were assigned randomly to each of the 3×3 between-participant cells of the design. The two orthogonal factors were the situation manipulation (i.e., object, setting, or event condition) and the catapult system (System 1, 2, or 3). Participants judged the function of all five objects from only one of the three catapult systems, with the five objects presented in a random order. For each object, participants described what they thought its function was. Later, blind judges coded these responses with respect to the a priori functional categories (i.e., fulcrum, lever, weight, projectile, target). Thus, the dependent variable was the rated accuracy of the inferred function. Additionally, participants' generated functions were coded to establish their situational content.

4.3.2. Materials

The three catapult systems used in Experiment 1 were also used here.

4.3.3. Procedure

Participants were tested individually. The general instructions stated that participants would see several objects all belonging to the same system, each of which served a particular goal in the overall system. Participants were also told that they would be asked to answer questions about each object. Participants were then taken to an adjoining room, and depending on the condition, shown different levels of situational information about a single system. Participants' responses were recorded with a digital audio recorder and later transcribed verbatim into a text file by a research assistant unaware of the experiment's purpose.

In the object condition, the experimenter took each component of a single system from a box and asked participants to describe its function. Objects were presented sequentially in a different random order for each participant, with only one object visible at a time. After describing the individual functions of the five objects, the experimenter placed all objects in a random arrangement on the table and the participant was asked to describe how the entire system worked.

In the setting condition, participants were presented with the system already assembled on top of a table. Participants did not see the system in use, but only statically. The experimenter pointed to each object sequentially in a random order and asked the participant to describe its function. The participant was then asked to describe how the entire system worked, with the system still in view.

In the event condition, participants viewed the assembled system and then watched the experimenter demonstrate it two times. In each demonstration, the experimenter placed the weight on the lever, which launched the projectile through the target. Following the two demonstrations, the experimenter pointed to each object in the static assembly in a random order and asked the participants to describe its function. The participant was then asked to describe how the entire system worked, with the system still in view.

4.4. Results

Participants often produced more than one response for a given object in a system. Excluding "don't know" responses, task-related commentaries, and response repetitions, the mean number of responses for a participant per object was 1.49 (min = 1, max = 4.2, SD = 0.74).

4.4.1. Accuracy rating and content coding

Two variables were scored for each response to an object in a system: (1) rated accuracy relative to the object's a priori function and (2) coded situational content. Each participant's description of the overall system function was also rated for accuracy relative to the system's a priori function. The first author performed the rating and coding judgments, in conjunction with a research assistant who established reliability. To prevent bias, information about each participant's experimental condition was removed from his or her data file prior to coding. Response repetitions and task-related commentaries were coded but not included in all analyses to follow.

For accuracy, each response was rated for how accurately it described the object's a priori function. Accuracy ratings were made on a three-point scale, with 0 indicating a completely inaccurate response, 1 indicating a partially accurate response, and 2 indicating an accurate response. The same scale was also used to rate the accuracy of each participant's description of how the complete system worked.

To code the content of participants' responses, the following coding categories were created: *physical structure*, *setting, object behavior, agent action*, and *agent goal* (see Table 1). When a response fell into two or more coding categories, each category was assigned to the relevant part of the response. For example,

Table 1
Examples of protocol statements for each content coding category in Experiment 2 from each level of situational information.

Content coding category	Situational information	System object	Coded response
Physical structure	Object	Projectile	It looks like a pyramid.
	Setting	Projectile	It has a tip.
	Event	Weight	The heavier part.
Setting	Structure	Projectile	Seems to be a base for some other object.
	Object	Weight	It is down.
	Event	Weight	It was on here [the lever].
Object behavior	Object	Lever	Can glide through the air.
	Setting	Projectile	Can slide.
	Event	Projectile	Moves through the air.
Agent action	Object	Fulcrum	You could put things on top.
	Setting	Weight	You could throw it.
	Event	Projectile	What you are shooting.
Agent goal	Object	Weight	You can use it to keep a door open.
	Setting	Lever	It is used to balance things.
	Event*	--	--

Note. Correct category and incorrect category were two additional content categories used to code participants' responses. The correct categories were *weight, projectile, lever, fulcrum,* and *target*. The incorrect categories were any other categories assigned to the objects.

* There were no goal responses in the event condition.

"you could throw it to hurt somebody" was segmented into "you could throw it" and "to hurt somebody," which were then coded as *agent action* and *goal*, respectively.

Frequently, participants attempted to place an object directly into an a priori category, rather than simply describing its function. When the category was correct, it was coded as a *correct category* (e.g., calling a projectile a "projectile"). When the category was incorrect, it was coded as an *incorrect category* (e.g., calling a projectile a "Christmas ornament").

To establish reliability, 15 participants (33% of the sample) were recoded independently by a research assistant who was blind to the hypothesis of the experiment, and to participants' conditions. Five participants were selected randomly from each level of the situation manipulation (i.e., object, setting, event), with the three catapult systems being approximately equal across levels. For content coding, the reliability for a given file was the total number of coincidences between the two judges divided by the total number of responses. Across the 15 participants, the average inter-rater agreement for content coding was 76.7%. Reliability for accuracy ratings was computed across all 76.7% of the responses for which the two judges agreed in how they coded their content (i.e., accuracy would not necessarily be expected to agree when the content codings differ). Across responses coded the same way for content, the Pearson correlation coefficient for the accuracy ratings between the two judges was 0.85.

4.4.2. Response accuracy for individual objects

Of primary interest was how well participants could infer the a priori function of an object from its object properties alone vs. with situational information. To answer this question, each participant's accuracy ratings for all five objects were averaged, and then submitted to a 3 (situation information) × 3 (system) between-participants ANOVA.

As Fig. 3 illustrates, situational information strongly affected accuracy ($F(2, 36) = 139.40$, $MSe = 0.07$, $p < 0.001$, $R^2 = 0.89$, power = 1). The object condition performed poorly (0.22 on the 0 to 2 accuracy scale), the setting condition showed some improvement (0.82), and the event condition performed near ceiling (1.80). As Fig. 3 illustrates further, the three different catapult systems differed marginally ($F(2, 36) = 3.09$, $MSe = 0.07$, $p < 0.058$, $R^2 = 0.15$, power = 0.56), and exhibited a marginal interaction with the amount of situational information ($F(4, 36) = 2.28$, $MSe = 0.07$, $p < 0.10$, $R^2 = 0.20$, power = 0.61). System 1 was easier to comprehend under setting instructions than the other two systems. Apparently, participants were better able to simulate the event for this system on the basis of setting information alone, without event information.

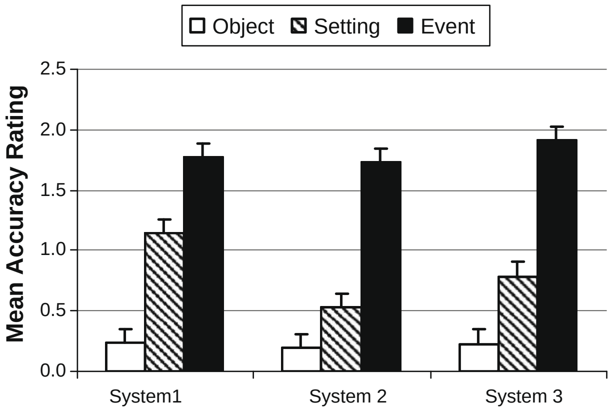

Fig. 3. Mean accuracy ratings in Experiment 2 as a function of situational information and system. Errors bars represent one standard error.

Planned comparisons were performed to further assess differences between levels of the situation manipulation. Setting information produced greater accuracy relative to object properties alone, $F(1, 36) = 38.96$, $MSe = 0.07$, $p < 0.001$. Event information further improved performance relative to the addition of setting information, $F(1, 36) = 105.91$, $MSe = 0.07$, $p < 0.001$.

To explore the marginal effect of system, all three systems were compared pairwise. Overall, responses to System 1 were most accurate (1.05), followed by responses to System 3 (0.98) and then to System 2 (0.82). The only significant difference was between System 1 and System 2 ($F(1, 36) = 5.93$, $MSe = 0.07$, $p < 0.05$). Systems 1 and 3 did not differ ($F(1, 36) = 0.61$, $MSe = 0.07$, $p < 0.50$), nor did Systems 2 and 3 ($F(1, 36) = 2.74$, $MSe = 0.07$, $p < 0.25$).

4.4.3. Response accuracy for the overall system

Accuracy ratings for the overall question about how the complete system worked were submitted to a 3 (situation information) × 3 (system) between-participants ANOVA. As in the analysis of individual object accuracy, situational information had a large effect ($F(2, 36) = 44.8$, $MSe = 0.22$, $p < 0.001$, $R^2 = 0.71$, power = 1). Accuracy was highest in the event condition (2.00), followed by the setting condition (0.93) and then the object condition (0.40). System had no effect ($F < 1$), and the two factors did not interact ($F < 1$).

Planned comparisons showed further that each additional level of situational information improved accuracy on the general question. Setting information produced greater accuracy relative to object properties alone, $F(1,\ 36) = 9.60$, $MSe = 0.22$, $p < 0.01$. Event information further improved performance relative to the addition of setting information, $F(1,\ 36) = 38.40$, $MSe = 0.22$, $p < 0.001$.

The same pattern of accuracy for individual object function and overall function suggests that the two measures were related. In fact, a Pearson correlation coefficient of 0.83 ($p < 0.001$) suggests that the situation manipulation affected accuracy at the object and system levels similarly.

4.4.4. Situational content

As described earlier, participants' responses were coded into categories for correct category, incorrect category, physical structure, setting, object behavior, agent action, and agent goal. Fig. 4A shows the average percentage of responses in each content category per participant across the five objects. In a 3 (situation information) × 7 (content) ANOVA performed on the arcsin transformation of these data, the situation manipulation showed a marginal effect ($F(2,\ 42) = 3.03$, $MSe = 0.003$, $p < 0.059$, $R^2 = 0.13$, power = 0.53), there was an effect of content category ($F(6,\ 252) = 18.61$, $MSe = 0.04$, $p < 0.001$, $R^2 = 0.31$, power = 1), and the two factors interacted significantly ($F(12,\ 252) = 18.08$, $MSe = 0.04$, $p < 0.001$, $R^2 = 0.46$, power = 1).

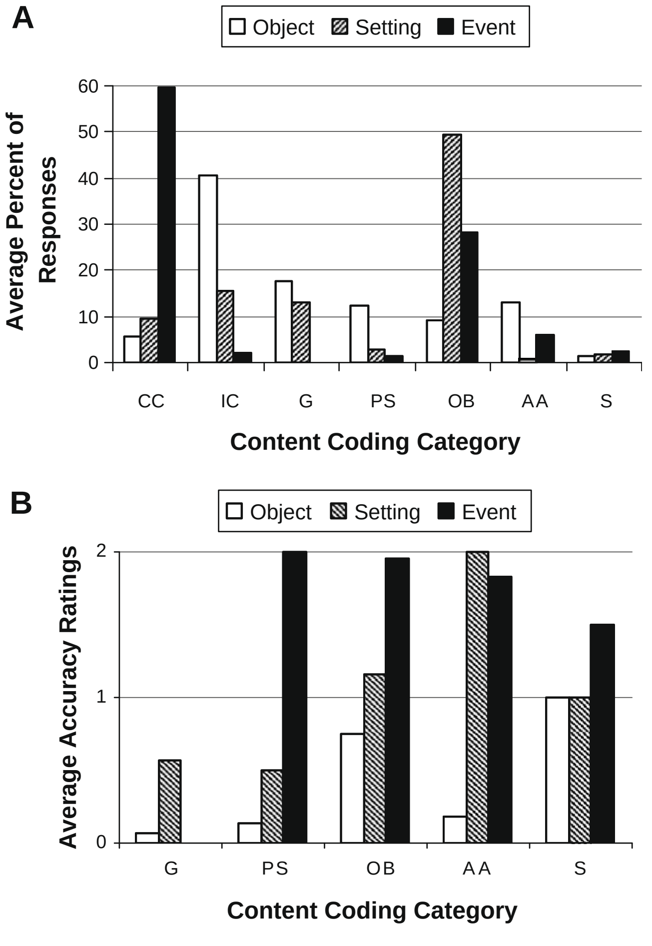

Fig. 4. Average percentage of responses (Panel A) and accuracy (Panel B) in Experiment 2 as a function of content coding category and situational information. CC is correct category, IC is incorrect category, G is goal, PS is physical structure, AA is agent action, OB is object behavior, S is setting. IC and CC are not included for the accuracy ratings in Panel B because, by definition, they were redundant with accuracy.

The average total frequency of responses across all five content categories was highest in the object condition (9.13), followed by the setting condition (6.53) and then the event condition (5.87). A post hoc interpretation of this result is that participants in the object condition had insufficient information and resorted to guessing, as they generated a series of low-confidence responses. In contrast, the setting condition had more information, allowing participants to settle into fewer, more confident interpretations. Finally, the event condition received enough situational information so that participants could describe objects' functions succinctly, rarely producing more than one response per object.

A related pattern in Fig. 4A concerns participants' production of categories. Although participants were asked to describe the function of each object, they frequently simply attempted to categorize the objects instead, as indicated by the high response rates for correct category (CC) and incorrect category (IC). Interestingly, categorization responses were the most frequent responses in the object and event conditions, probably for different reasons. In the object condition, the physical information presented was so uninformative that participants may not have been able to imagine functional situations in which the object could be used, and instead guessed categories. In the event condition, the physical, setting, and event information presented may have been so informative as to converge on the correct category. Unlike the other two conditions, the setting condition mostly produced responses that were not categorizations of the objects, but instead described situational information. Participants in this condition had enough information about object properties and the setting that they could intelligently infer non-presented information about the situation having to do with goals (G), object behavior (OB), and agent actions (AA). Nevertheless, they did not have enough information to produce high rates of correct categorization (CC).

Three further patterns in Fig. 4A are of interest. First, the object and setting conditions produced more goals (G) than the event condition, suggesting that these participants were trying to figure out what the goal of using an object might be. In contrast, the event condition produced virtually no goals, perhaps because the goals were so clear as not to be worth mentioning explicitly. Second, the setting and event conditions produced more object behaviors (OBs) than did the object condition (especially the setting condition). This suggests that the setting condition focused on possible behaviors that objects in the static system assembly might display if put in motion. Indeed, inferring object behaviors was the most common way in which the setting condition described object functions. Third, the object condition produced more descriptions of physical structure (PS) than the other two conditions, similar to the focus on object properties in our control experiment. Because information about each object in isolation was so impoverished, participants in the object condition often appeared to conceptualize objects in terms of their physical properties, not their situational ones.

Fig. 4B further demonstrates that increasing situational information presented for an object improved categorization. This figure shows the average accuracy rating for each type of situational content as a function of the situation manipulation (accuracy is not shown for the correct and incorrect category codes, because it would be redundant with the accuracy information in Fig. 4A). As can be seen, the general trend for each content category was for accuracy to be lowest in the setting condition, higher in the setting condition, and highest in the event condition. The one exception is for agent actions, where performance was at ceiling for both the setting and the event conditions. Because only one agent action was produced in the entire setting condition (Fig. 4A), this particular comparison should be interpreted with caution.

Most importantly, the general pattern of increasing accuracy in Fig. 3 for functional inferences across the situation manipulation also held across the specific content categories in Fig. 4. The more situational information presented for an object, the more likely it was to activate the correct category in memory, such that accurate information about the object was produced.

4.5. Discussion

An object's function is not obvious from its physical properties alone. It cannot be immediately inferred as an affordance (Gibson, 1979). Instead, perceiving the setting and events associated with the object's function facilitates the activation of category knowledge by matching this information in the category's representation in memory. Inferences about an object's function did not increase discretely, shifting from no understanding to complete understanding. Instead, accuracy increased in a graded manner. As increasing amounts of situational information became available, the likelihood of accessing the a priori categories increased, which then made the requested information about function available.

Interestingly, the most common content produced in the event condition was the correct category for the object (i.e., *fulcrum, lever, weight, projectile*, or *target*). This strongly suggests that setting and event information constitute core properties of these categories' representations in memory.

Participants in the object and the setting conditions also generated functions to some extent, just not the functions associated with the a priori categories. Specifically, the large number of incorrect categories, goal statements, and action statements that these participants produced suggests that they were trying to infer functions. More importantly, these participants often appeared to be *guessing* possible functions. Because no particular function of the object was obvious in its impoverished context, participants tried to infer setting and event contexts that might suggest one. Most importantly, however, the a priori categories were not obvious without setting and event information. Until this information was provided, these categories were not salient.

5. Experiment 3

This final experiment further assessed the impact of situational information on inference. Rather than requesting specific inferences from participants, as in Experiment 2, we allowed participants to draw inferences in a more open-ended manner. On each trial, participants were shown an object and asked to generate properties for "objects of this type."

As in the previous experiments, the amount of situational information presented varied across the object, setting, and event conditions. If people represent object categories with situational information, then the more of this information presented, the more likely it is that the a priori categories will be accessed. If the correct category is accessed, then the properties that participants generate should be more appropriate. Two specific measures of property appropriateness were assessed: property coverage and relational properties. Each is addressed in turn.

5.1. Property coverage

This measure assesses how well each property generated for a category covers the set of entities that, in principle, could instantiate the category (cf. Osherson, Smith, Wilkie, Lopez, & Shafir, 1990). If a generated property covered all or most category members in an informative and not too general manner, it received a coverage rating of 1. Imagine that a participant generated *is propelled* as a property of a projectile object. Because this property is true of all projectiles, and because it is informative about category members, it would receive a coverage rating of 1. Alternatively, a generated property received an intermediate coverage score of 0.5 if it was true of all category members, but was so general as not to be informative. For example, *can hold objects* is generally true of most fulcrum objects, but is so general that it describes many other objects as well (e.g., shelf, table, box), so as not to be highly informative about the fulcrums. Finally, a coverage score of 0 was given to a property that only covered a minimal subset of category members (e.g., *pyramidal* for a fulcrum), or that was irrelevant to the function (e.g., *has wire on the top* for a weight). Often these latter properties only described the object currently serving as a cue for generating properties.

Property coverage was used to assess the hypothesis that situational information is stored with categories and plays a central role in categorization. If this hypothesis is correct, then as more situational information accompanies an object during categorization, the likelihood of correct categorization should increase. As the likelihood of correct categorization increases, the properties that become active for the category should become increasingly accurate descriptions of it. Coverage should become increasingly appropriate, approaching 1 on the coverage scale.

5.2. Relational properties

The second measure of property appropriateness assessed in Experiment 3 was the generation of relational physical properties. Situations are inherently relational, typically containing predictable configurations of objects, agents, and events. A wide variety of spatial, temporal, causal, and intentional relations underlie these configurations. If situational information underlies the representation of objects, then as this information becomes active, it should activate the relations that organize it.

Of particular interest here are relations that are physical properties of objects, what Barr and Caplan (1987) referred to as *extrinsic properties* (as opposed to intrinsic properties). For example, the correct size of a toothbrush for achieving its function is not an intrinsic property of an object per se, but depends extrinsically on the settings and events in which it will be used (e.g., brushing the mouth of a cat vs. an elephant). Conversely, the shape of a toothbrush is an intrinsic property, because it does not depend on settings or events. Previous research had shown that relational physical properties play important roles when people learn functional object categories like those studied here. For example, the relative length and rigidity of an artifact were important for the categories in Kemler-Nelson (1999). Relative length was also important for the categories in Hauser (1997).

In Experiment 3, we predicted that the number of relational physical properties in participants' property listing protocols would increase with the amount of situational information presented. To the extent that situational information converges on the correct category in memory, accurate situational knowledge about the category should become active. Because relational physical properties are central to organizing this knowledge, they in turn should become active, such that participants produce them. In contrast, when situational information is not present, and the correct category does not become active along with related situational knowledge, relational physical properties should be produced less often.

5.3. Experiment overview

To assess property coverage and relational physical properties, we again manipulated the amount of situational information that participants received. As in the previous experiment, three levels of situational information were manipulated between participants: object, setting, and event. In the object condition, participants received the five objects from a system in a randomized order. In the setting condition, participants viewed an assembled system but did not see it operate. In the event condition, participants watched demonstrations of an assembled system. In all three conditions, participants were simply asked to describe the properties of each object from the system that they viewed. We then assessed the differential effects of increasing situational information on property coverage and on relational properties.

To keep the analysis of property listing data manageable, only two systems were included in this experiment. Specifically, we included what appeared to be the easiest system (System 1) and the hardest system (System 2), based on the previous experiments. Thus, the two systems that were used spanned the range of difficulty.

Experiment 3 assessed an additional hypothesis of interest: Do the benefits of situational information for conceptualizing object categories increase after viewing multiple category instances as opposed to viewing just one instance? To assess this hypothesis, an additional group of participants received *both* the easy and the hard systems before producing properties, rather than just receiving one system. In the object condition, participants viewed the corresponding functional objects from the two systems in isolation. In the setting condition, participants viewed both assembled systems. In the event condition, participants watched demonstrations of both systems. All participants then generated properties for each of the five pairs of corresponding objects. Of interest was whether viewing two objects from an object category produced higher property coverage and relational properties than just viewing one object.

5.4. Method

5.4.1. Participants and design

Participants were 108 Emory undergraduates who participated for course credit (99 females, 9 males). The experiment followed a 3×3 between-participants design, with three levels of situational information (object, setting, event) and three types of system (easy, hard, both easy and hard).

5.4.2. Materials and procedure

Only Systems 1 and 2 from the previous experiments were used. Based on the results of previous experiments, System 1 was designated the "easy system," and System 2 was designated the "hard system." In all conditions, the system(s) whose objects were to be described remained covered during the instruction period, and were only uncovered when participants were ready to produce properties. Participants were allowed to produce properties until no more came to mind, typically about 20 seconds per object.

In the object condition, the objects were presented individually one at a time in a random order. No other objects from the system were visible. While an object was in view, participants wrote down a list of typical properties for that "type of object." The object was then hidden from view, and a new object presented. When participants received both systems, the two objects sharing the same a priori function were presented simultaneously, and participants were asked to describe that type of object.

In the setting condition, the experimenter pointed sequentially to each object in its assembled system in a random order. Participants produced a list of properties typical for that type of object, with this process continuing until all objects had been described. When participants received both systems, the two assembled systems were shown simultaneously, the experimenter pointed to the analogous objects sharing the same a priori function in a random order, and participants were asked to describe that type of object.

In the event condition, the experimenter performed two consecutive demonstrations with the system. The experimenter then pointed sequentially to each object in the assembled system in a random order and asked participants to describe the properties typical for each type of object. When participants received both systems, the two assembled systems were demonstrated sequentially and in a random order. The experimenter then pointed to analogous objects sharing the same a priori function in the assembled system in a random order, and participants were asked to describe that type of object.

5.5. Results

Participants produced an average of 4.06 properties for each of the 5 objects ($SD = 1.59$). The average number of properties produced did not vary as a function of situational information ($F < 1$), or system ($F(2, 99) = 1.21$, $MSe = 62.5$, $p < 0.5$, $R^2 = 0.02$, power = 0.26), nor was there an interaction ($F(4, 99) = 1.52$, $MSe = 62.5$, $p < 0.25$, $R^2 = 0.06$, power = 0.45).

5.5.1. Property content

The properties that each participant produced for an object were coded for their conceptual content by the first author, blind to participants' conditions. The coding categories were motivated by the HIPE theory of function (Barsalou et al., 2004), and are shown in the first column of Table 2. These coding categories differed somewhat from those in earlier experiments, given that participants produced a much broader range of properties here that required a broader set of coding options. These categories included three types of *history* codes (an object's role, its invention, and its manufacturing); two types of *property* codes (independent object properties and relational object properties); three types of *setting* codes (other objects in the system besides the focal object, the agent's goal, and the agent's action);[2] and three types of *event* codes (focal object behavior, behavior of other objects in the setting, outcome). A variety of *other* content codes captured related *categories* (e.g., "it's a table"); *don't know responses*; responses that were potentially *relevant* but that were not captured by the coding scheme (i.e., properties relevant to the object's function that did not fit into an existing coding category, e.g., "it depends on other objects to function"); and *irrelevant* responses (i.e., properties that were irrelevant to the object's function, e.g., "versatile", "simple"). *Other relevant* responses occurred less than 1% of the time, demonstrating that the primary coding categories came close to exhausting the relevant content produced. A research assistant blind to the hypothesis and participants' conditions independently coded all projectile and weight properties (approximately 40% of the total number of properties produced across all participants). Inter-coder reliability was 80% for the projectile properties, and 71% for the weight properties.

The average proportion for each type of content coded for each participant was computed across the five types of functional objects. The columns to the right in Table 2 present the average proportions across participants in the different situation conditions.

5.5.2. Property coverage

The first author rated each property generated for an object by a participant for coverage, again blind to participants' conditions. As described earlier, a scale of 0, 0.5, and 1 was used. Reliability was assessed by having a research assistant code all properties generated for the projectiles and weights, again blind to the hypotheses and participants' conditions. Inter-rater reliability was $r = 0.81$ for the projectile coverage scores and $r = 0.80$ for the weight coverage scores.

The average coverage score for each participant across all generated properties was submitted to a 3 (situational information) $\times$ 3 (system) between-participants ANOVA. As Fig. 5 illustrates, sit-

---

[2] In HIPE, agent goal and agent action are setting properties because they belong to the initial setting that triggers the subsequent event sequence that achieves an object's function (see Barsalou et al., 2004, for details).

Table 2
Percentage of each content coding category as a function of situational information in Experiment 3.

Content domain	Content category	Situational information		
		Object	Setting	Event
History	Role	15.3	12.3	13.7
	Invention	0.1	0.1	0.1
	Manufacturing	4.8	2.2	1.8
Property	Independent	43.0	41.6	36.4
	Relational	5.9	8.4	13.0
Setting	Other object	1.5	5.8	4.4
	Agent goal	0.1	0	0.4
	Agent action	8.9	8.0	2.1
Event	Focal object behavior	4.7	6.9	8.6
	Other object behavior	0.3	2.1	1.9
	Outcome	0.5	0	0.3
Other	Category	9.9	8.9	12.9
	Don't know	1.9	1.1	0.5
	Relevant	0.1	0.6	0.8
	Irrelevant	2.8	1.9	3.0

Note. Total number of responses produced: object = 744, setting = 722, event = 729.

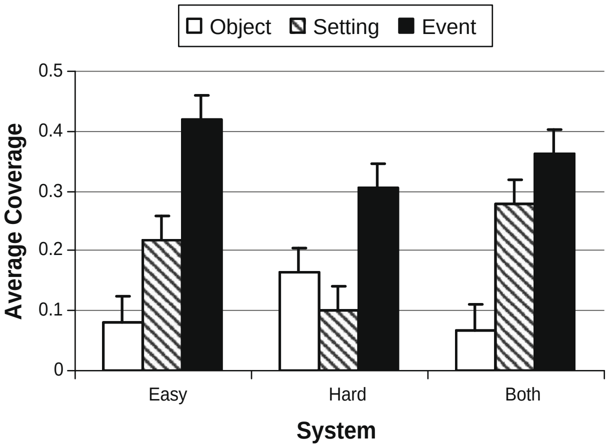

Fig. 5. Average coverage in Experiment 3, as a function of situational information and system (easy, hard, or both). Error bars represent one standard error.

uational information had the predicted effect. As setting and event information became increasingly available during categorization, the properties that participants produced exhibited increasingly optimal coverage of the functional categories ($F(2, 99) = 28.73$, $MSe = 0.021$, $p < 0.001$, $R^2 = 0.37$, power = 1). System had no overall effect ($F(2, 99) = 1.29$, $MSe = 0.021$, $p < 0.50$, $R^2 = 0.03$, power = 0.27), but interacted significantly with situational information ($F(4, 99) = 3.34$, $MSe = 0.021$, $p < 0.05$, $R^2 = 0.12$, power = 0.83).

To assess the source of the interaction, further comparisons were performed. For the hard system, event information produced greater coverage than setting information ($F(1, 99) = 11.75$, $MSe = 0.021$, $p < 0.001$), but setting information did not produce greater coverage than object information ($F(1, 99) = 1.16$, $MSe = 0.021$, $p < 0.50$). This indicates that setting information was not sufficient for activating the a priori categories in the hard system.

In contrast, setting information for the easy system did produce greater coverage than did object information ($F(1, 99) = 5.2$, $MSe = 0.021$, $p < 0.05$), indicating that participants could infer the a priori categories to some extent from setting information. Event information further improved the access of the a priori categories relative to setting information for the easy system ($F(1, 99) = 11.56$, $MSe = 0.021$, $p < 0.001$).

Finally, the two-systems condition showed still a different pattern. When participants received both systems, setting information produced a particularly large increase in coverage relative to the object condition ($F(1, 99) = 12.33$, $MSe = 0.021$, $p < 0.001$). This increase was sufficiently large that adding event information did not produce significant improvement ($F(1, 99) = 2.04$, $MSe = 0.021$, $p < 0.25$). Perceiving two static systems may have allowed participants to simulate events of the systems in operation, such that sufficient situational information was available to access the a priori categories.

When event information was available, however, viewing two systems did not improve coverage relative to viewing one ($F < 1$). This suggests that receiving one instance of complete situational information—a setting and an event—was sufficient to access the a priori categories. Viewing an additional instance provided no further benefit.

5.5.3. Relational physical properties

Relational physical properties of the focal object were defined as those that depended, in part, on something else in the setting (e.g., a projectile is "light" relative to its counterweight which is "heavy", a fulcrum is "stable" to allow its lever to "pivot" on it). As Table 2 illustrates, the proportion of relational physical properties more than doubled across the object (0.06), setting (0.08), and event (0.13) conditions. A 3 (situational information) × 3 (system) between-participants ANOVA on the arcsin transformation of these proportions confirmed this pattern ($F(2, 99) = 8.8$, $MSe = 0.006$, $p < 0.001$, $R^2 = 0.15$, power = 0.97). Pairwise comparisons further indicated that receiving event information increased relational physical properties relative to setting information ($F(1, 99) = 9.58$, $MSe = 0.006$, $p < 0.01$), but that receiving setting information did not produce an increase relative to object information ($F < 1$). System had no effect ($F < 1$). There was no interaction ($F < 1$).[3]

The lack of both a system effect and an interaction is interesting. As Fig. 6 illustrates, receiving the easy system, the hard system, or both systems had no effect on the production of relational physical properties. Most notably, when participants in the events condition observed both the easy and the hard systems together, they did not generate a higher proportion of relational physical properties (0.12) than when they received the easy system alone (0.15) or the hard system alone (0.13). This suggests that observing a single system with all the relevant situational information was sufficient to produce deep understanding of the relational structure that underlies a functional object category.

5.6. Discussion

As situational information increased, the properties that participants produced became increasingly accurate. Specifically, these properties became increasingly optimal in coverage, not being too general, too specific, or irrelevant. When situational information was not present, participants were much less likely to activate property information for a category that covered its exemplars adequately. Similarly, the proportion of relational physical properties increased as more situational information became available. When the correct categories became active, they provided access to relational properties that structured the perceived situations.

Receiving two systems as opposed to one did not improve object descriptions in the event condition. When participants only viewed a single system in operation, property coverage was as accurate as when they viewed two systems. Similarly, the production

---

[3] As can be seen in Table 2, the proportions for independent physical properties (e.g., color, shape) decreased as increasing amounts of setting and event information were available, suggesting that these properties became less salient as more relational object properties became available. However, an ANOVA on the proportions for the independent object properties showed no effect, largely because of high variance between participants.

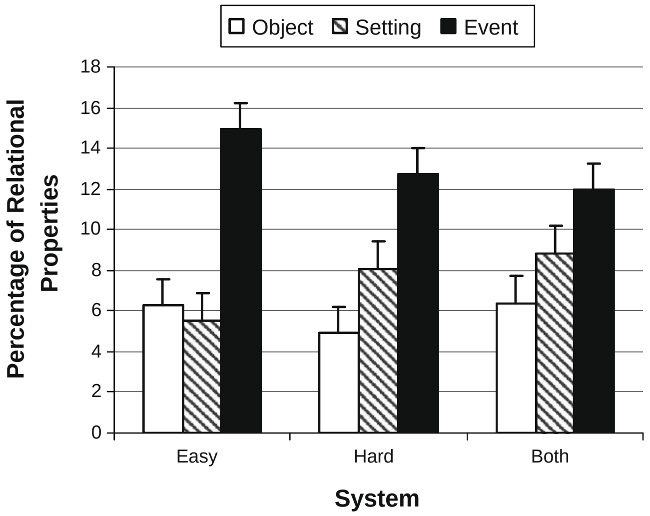

Fig. 6. Percent of relational properties in Experiment 3 as a function of situational information and system (easy, hard, or both). Error bars represent one standard error.

of relational physical properties was comparable when viewing one system vs. two systems. These findings indicate that the presence of adequate situational information for one system was sufficient to access the a priori categories, such that the information produced for them was near optimal.

Viewing two systems in the setting condition, however, showed a different pattern. When participants viewed two assembled systems, but did not see them operate, they appeared to infer the events associated with system operation. As a result, adding event information did not produce further benefit. Inducing relevant events from static assembled systems appears to benefit from viewing multiple systems, rather than only viewing one.

6. General discussion

Three experiments demonstrated that situational information plays central roles in object categorization and inference. When no situational information was present, participants primarily viewed objects in terms of their physical properties, which were not sufficient to activate the a priori functional object categories. When situational information was present, participants increasingly viewed objects in terms of their relational physical properties (Experiment 3), and activated the a priori functional categories at much higher rates (across all experiments). As evidenced by the content of the properties that participants generated, situational information promoted changes in the pattern of inferences, in manners consistent with the greater incidence of situational information (Experiments 2 and 3).

Though increasing situational information improved accuracy in a graded manner, event information had to be present to maximally match category representations in memory. The presence of setting information produced some benefit, but not nearly as much as setting and event information together (Experiments 2 and 3). Thus, setting information resides in the representations of these object categories, but its presence during categorization is often not sufficient to activate them. Instead, event information must also be present to maximally match category representations in memory.

Situational information affected both the initial access of categories (categorization) and their subsequent use (inference). Experiment 1 used clustering to measure categorization. When no situational information was provided, the a priori functional categories were seldom accessed and thus did not organize the objects. When complete situational information was provided, the a priori functional categories were accessed at high rates and produced corresponding clusters. Although Experiments 2 and 3 focused on inference following categorization, they offered indirect evidence that categorization benefited from situational information. Because the inferences drawn about an object depended on the category accessed initially, accurate inferences reflected accurate categorization. Experiment 2 showed that the accuracy of functional inferences increased with the amount of situational information present during categorization. Experiment 3 showed that inferences about many kinds of properties became more accurate as situational information increased.

6.1. Property familiarity

6.1.1. Familiar object properties

The experiments reported here used novel objects that participants had never seen before. Although participants readily assigned these objects to the a priori functional categories when relevant situational information was available, they could not do so when situational information was not available. Using novel objects allowed us to demonstrate the presence of situational knowledge in memory, along with its central roles in categorization and inference.

Clearly, however, people can categorize some isolated objects correctly in the absence of situational information, as many experiments show. When participants see pictures of typical hammers, chairs, and hats, for example, they categorize them correctly in the absence of situational information. Such findings demonstrate that object information alone can be sufficient to activate correct categories. When familiar diagnostic properties are perceived for an object, they carry enough information to activate the correct category in memory. It does not follow, however, that situational information is irrelevant to categorization. Research on object recognition has shown that the presence of relevant setting information speeds the categorization of familiar category members recognizable solely on the basis of their object properties (e.g., Bar & Ullman, 1996; Biederman, 1981; Dobel et al., 2007; Palmer, 1975). Research has similarly shown that motion events speed the categorization of familiar category members (e.g., Stone, 1999; Vuong & Tarr, 2004).

Situational information is also central during categorical inference. In many experiments, when participants receive familiar objects during categorization, they nevertheless activate extensive situational information as inferences. For example, many studies have presented participants with pictures of familiar category members and found that motor (event) inferences became active. Behavioral studies have demonstrated such inferences (e.g., Bub & Masson, 2006; Glover, Rosenbaum, Graham, & Dixon, 2004; Riggio et al., 2008; Tucker & Ellis, 1998; Tucker & Ellis, 2004), as have neuroimaging studies (e.g., Chao & Martin, 2000; Creem-Regehr & Leeb, 2005; Kellenbach, Brett, & Patterson, 2003). Many studies have similarly demonstrated that the isolated pictures of objects activate setting information (e.g., Aminoff, Gronau, & Bar, 2007; Bar, 2004), as do words for objects (e.g., Vallée-Tourangeau, Anthony, & Austin, 1998).

All this research is consistent with the conclusion that object properties are not the sole factor in categorization and inference, even when these properties are familiar. To the contrary, situational information plays central roles. Indeed, inferred situational information may be the point of conceptualization, as it prepares agents for situated action with a category (e.g., Barsalou, 2003b; 2005b; 2009).

6.1.2. Familiar setting properties

The presence of a familiar setting can facilitate categorization and inference. In Experiments 2 and 3, System 1 exhibited higher accuracy than the other systems. When participants saw the objects in System 1 assembled but static, they were often able to infer the correct categories. They did not need to see events to infer them. This suggests that the configured objects in System 1 matched setting information in memory well enough to activate the categories.

As Panel B of Fig. 1 illustrates, the fulcrum and lever in System 1 best approximated a typical fulcrum and lever. The fulcrums and levers in Systems 2 and 3 were much less typical. This familiar object configuration in System 1 may have been sufficient to access the categories of *fulcrum* and *lever*. Once these categories became active, they may have activated event simulations, which in turn helped to categorize other objects as *weight, projectile,* and *target*.

Like object properties, setting properties vary in familiarity. As the familiarity of both increases, their ability to activate event simulations increases, thereby influencing categorization and inference.

6.1.3. Familiar event properties

The events in our experiments—together with the settings and objects—were usually sufficient to activate the a priori functional categories. This suggests that all three sources of information together were necessary and sufficient for accurate categorization. Alternatively, the events may have been highly recognizable on their own if they were familiar to participants (unlike the object and setting properties).

One possibility is that events are represented in a fundamentally different manner than objects and settings. Whereas objects and setting are represented in brain systems that process physical entities and the relations between them, events and actions are processed in brain systems that process motor and motion information.[4] Furthermore, these two systems may vary considerably in the abstractness and generality of their representations. Whereas objects and settings may be represented in relatively precise and narrow formats, actions and events may be represented in relatively broad and qualitative formats. As a result, novel objects and settings may usually be recognized only when they are relatively close to familiar objects and settings in memory. Conversely, novel actions and events may often be recognized even when relatively different from the previously experienced events and actions. Although the object and setting properties of the novel objects here may have appeared unfamiliar to participants, the event properties may have appeared familiar.

Thus, our experiments do not distinguish between the possibility that all three forms of information—object, setting, and event—were sufficient to access the functional object categories, or whether event information alone was sufficient. Regardless, it is clear that situational information plays major roles in categorization and inference.

6.1.4. The time course of object, setting, and event information in categorization

Another important issue concerns whether these information sources provide staged access to categories, with one source providing initial access, and the others providing later access. Perhaps object properties access functional object categories first, with setting and event properties operating later.

We suspect that this may be an artificial issue. In every day life, people continually perceive settings that contain objects and events in a highly dynamic flux (Barsalou, Breazeal, & Smith, 2007). Furthermore, people often anticipate objects and events that are likely to occur in the near future. The categorization of functional objects occurs against this backdrop. Representations of the setting and current events are likely to prime objects that are likely to be observed, such that situational information—not object properties—provides initial access to the respective object categories. Only when an object is unexpected may object properties provide initial access. Otherwise, it is likely that many forms of information, in constant flux, contribute to the activation of anticipated categories.

For these reasons, we believe that much is to be learned from studying the flux of information that occurs during situated action. We doubt that studying categorization in the standard manner of presenting limited snapshots of object, setting, and event information in rigid temporal sequences is likely to shed much light on categorization in these more complicated situations. At a minimum, it is necessary to better understand the structure and dynamics of information in situations before committing to the idealized laboratory paradigms for studying specific processes within them.

6.1.5. Grounding situational information in the simulation mechanism

How does situational information contribute to object categorization? Our results are consistent with the following view. As people experience situations that contain the same focal object, they store sensory-motor memories of the object together with information about the surrounding situations, including settings and events. On later occasions, when an instance of the focal object category occurs in isolation and its perceptual representation matches representations of the object category in memory, memories of associated situations become active and run as simulations that go beyond the information perceived for the object. Thus, these simulations of situations constitute a mechanism for prediction and inference, providing information about likely settings and events.

As the amount of situational information perceived for an object category increases, categorization of the object becomes increasingly accurate. When no situational information is present, category access is unreliable, and simulations for incorrect categories may fire, making categorization and inference relatively inaccurate (as in the object condition across experiments). When, however, a setting is present, or a setting and an event, the increasing amount of situational information overlaps simulations associated with the object in memory, increasing the likelihood that correct categorizations are made (as in the setting and event conditions of the experiments here).

When an object has diagnostic object properties familiar to the perceiver, categorization is likely to be accurate even in isolation, given that situational information is not necessary for distinguishing between categories in memory. Nevertheless, predictive simulations should still follow as inferences (e.g., perceiving a typical hammer and simulating the grasping of it). Descriptions of this view and its extensions to situational information can be found in Barsalou, 1999a,b; Barsalou, 2003a,b; Barsalou, 2005a,b; Barsalou, 2008; Barsalou, 2009), Barsalou et al. (2003), and Yeh and Barsalou (2006).

Like the objects in Fig. 1, functional object categories (tools, utensils, artifacts) vary widely in their object properties (e.g., different types of can openers). Thus, object properties may not be sufficient for categorization, and simulations of surrounding situations have greater potential for contributing to both categorization and inference (as in the event condition across experiments). As discussed earlier, the fact that perceiving an event was necessary for participants to achieve reliable categorization and accurate inferences demonstrates that these events are stored in memory with functional object categories, and that this information plays central roles during categorization and inference.

---

[4] Our design does not allow us to decide which, motor or motion information, was more important for our materials. However, content produced in the property listing tasks suggests that motion dominated. In Experiments 2 and 3, motor related content (i.e., agent actions) was stably low, or decreased as more situational information was provided. In contrast, motion related content (i.e., focal object behavior) increased (see Table 2 and Fig. 4).

Acknowledgements

We are grateful to Shurin Hase, Courtney Emery, and Joy Lynn Brasfield for assistance in running these experiments. This work was supported by National Science Foundation Grants SBR-9421326, SBR-9796200, and BCS-0212134 to Lawrence Barsalou, and by financial support from the University of Tarapaca to Sergio Chaigneau. Address correspondence to Sergio Chaigneau, Universidad Adolfo Ibáñez, Escuela de Psicología, Avenida Diagonal Las Torres 2640, Peñalolén, Santiago-Chile (sergio.chaigneau@uai.cl), or to Lawrence W. Barsalou, Department of Psychology, Emory University, Atlanta, GA 30322 (barsalou@emory.edu, http://userwww.service.emory.edu/~barsalou/.).

References

Aminoff, E., Gronau, N., & Bar, M. (2007). The parahippocampal cortex mediates spatial and non-spatial associations. *Cerebral Cortex, 17*, 1493–1503.

Bar, M. (2004). Visual objects in context. *Nature Reviews Neuroscience, 5*, 617–629.

Bar, M., & Ullman, S. (1996). Spatial context in recognition. *Perception, 25*, 343–352.

Barr, R. A., & Caplan, L. J. (1987). Category representations and their implications for category structure. *Memory & Cognition, 15*, 397–418.

Barsalou, L. W. (1999a). Language comprehension: Archival memory or preparation for situated action. *Discourse Processes, 28*, 61–80.

Barsalou, L. W. (1999b). Perceptual symbol systems. *Behavioral and Brain Sciences, 22*, 577–609.

Barsalou, L. W. (2003a). Abstraction in perceptual symbol systems. *Philosophical Transactions of the Royal Society of London: Biological Sciences, 358*, 1177–1187.

Barsalou, L. W. (2003b). Situated simulation in the human conceptual system. *Language and Cognitive Processes, 18*, 513–562.

Barsalou, L. W. (2005a). Abstraction as dynamic interpretation in perceptual symbol systems. In L. Gershkoff-Stowe & D. Rakison (Eds.). *Building object categories. Carnegie Symposium Series* (pp. 389–431). Mahwah, NJ: Erlbaum.

Barsalou, L. W. (2005b). Continuity of the conceptual system across species. *Trends in Cognitive Sciences, 9*, 309–311.

Barsalou, L. W. (2008). Grounded cognition. *Annual Review of Psychology*, 59.

Barsalou, L.W. (2009). Situating concepts. In P. Robbins M. Aydede (Eds.), Cambridge handbook of situated cognition. New York: Cambridge University Press.

Barsalou, L. W., Breazeal, C., & Smith, L. B. (2007). Cognition as coordinated non-cognition. *Cognitive Processing, 8*, 79–91.

Barsalou, L. W., Niedenthal, P. M., Barbey, A., & Ruppert, J. (2003). Social embodiment. In B. Ross (Ed.). *The Psychology of Learning and Motivation* (Vol. 43, pp. 43–92). San Diego: Academic Press.

Barsalou, L. W., Sloman, S. A., & Chaigneau, S. E. (2004). The HIPE theory of function. In L. Carlson & E. van der Zee (Eds.), *Representing functional features for language and space. Insights from perception, categorization and development* (pp. 131–147). Oxford: Oxford University Press.

Barwise, J., & Perry, J. (1983). *Situations and attitudes*. Cambridge, MA: MIT Press.

Biederman, I. (1972). Perceiving real-world scenes. *Science, 177*, 77–80.

Biederman, I. (1981). On the semantics of a glance at a scene. In M. Kubovy & J. R. Pomerantz (Eds.), *Perceptual Organization* (pp. 213–253). Hillsdale, NJ: Erlbaum.

Bransford, J. D., & Johnson, M. K. (1973). Considerations of some problems of comprehension. In W. G. Chase (Ed.), *Visual information processing*. New York: Academic Press.

Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence, 47*, 139–159.

Bub, D. N., & Masson, M. E. J. (2006). Gestural knowledge evoked by objects as part of conceptual representations. *Aphasiology, 20*, 1112–1124.

Chaigneau, S. E., Barsalou, L. W., & Sloman, S. (2004). Assessing the causal structure of function. *Journal of Experimental Psychology: General, 133*, 601–625.

Chao, L. L., & Martin, A. (2000). Representation of manipulable man-made objects in the dorsal stream. *Neuroimage, 12*, 478–484.

Cheng, P. W., & Holyoak, K. J. (1985). Pragmatic reasoning schemas. *Cognitive Psychology, 17*, 391–416.

Clark, A. (1997). *Being there: Putting brain, body, and world together again*. Cambridge, MA: MIT Press.

Clark, H. H. (1992). *Arenas of language use*. Chicago: University of Chicago Press.

Creem-Regehr, S. H., & Leeb, J. N. (2005). Neural representations of graspable objects: Are tools special? *Cognitive Brain Research, 22*, 457–469.

Dobel, C., Gumnior, H., Bölte, J., & Zwitserlood, P. (2007). Describing scenes hardly seen. *Acta Psychologica, 125*, 129–143.

Dunbar, G. (1991). *The cognitive lexicon*. Tübingen: Gunter Narr Verlag.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton-Mifflin.

Gick, M. L., & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology, 12*, 306–355.

Glenberg, A. M. (1997). What memory is for. *Behavioral and Brain Sciences, 20*, 1–55.

Glover, S., Rosenbaum, D. A., Graham, J. R., & Dixon, P. (2004). Grasping the meaning of words. *Experimental Brain Research, 154*, 103–108.

Goldberg, A. E. (1995). *Constructions: A construction grammar approach to argument structure*. Chicago: University of Chicago Press.

Greeno, J. G. (1998). The situativity of knowing, learning, and research. *American Psychologist, 53*, 5–26.

Hauser, M. (1997). Artifactual kinds and functional design features: what a primate understands without language. *Cognition, 64*(3), 285–308.

Hauser, M., Pearson, H., & Seelig, D. (2002). Ontogeny of tool use in cottontop tamarins, Saguinus oedipus: Innate recognition of functionally relevant features. *Animal Behaviour, 64*(2), 299–311.

Johnson-Laird, P. N. (1983). *Mental models*. Cambridge, MA: Harvard University Press.

Kellenbach, M. L., Brett, M., & Patterson, K. (2003). Actions speak louder than functions: The importance of manipulability and action in tool representation. *Journal of Cognitive Neuroscience, 15*, 30–46.

Kemler-Nelson, D. G. (1995). Principle-based inferences in young children's categorization: Revisiting the impact of function on the naming of artifacts. *Cognitive Development, 10*, 347–380.

Kemler-Nelson, D. G. (1999). Attention to functional properties in toddlers' naming and problem-solving. *Cognitive Development, 14*, 77–100.

Kemler-Nelson, D. G., Russell, R., Duke, N., & Jones, K. (2000). Two-year-olds will name artifacts by their functions. *Child Development, 71*(5), 1271–1288.

Kirsh, D. (1991). Today the earwig, tomorrow man. *Artificial Intelligence, 47*, 161–184.

Madole, K. L., & Oakes, L. M. (2004). Infants' attention to and use of functional properties in categorization. In L. Carlson & E. van der Zee (Eds.), *Representing functional features for language and space: Insights from perception, categorization and development*. Oxford: Oxford University Press.

Mischel, W. (1968). *Personality and assessment*. New York: Wiley.

Osherson, D. N., Smith, E. E., Wilkie, O., Lopez, A., & Shafir, E. (1990). Category based induction. *Psychological Review, 97*, 185–200.

Palmer, S. (1975). The effects of contextual scenes on the identification of objects. *Memory & Cognition, 3*, 519–526.

Riggio, L., Iani, C., Gherri, E., Benatti, F., Rubichi, S., & Nicoletti, R. (2008). The role of attention in the occurrence of the affordance effect. *Acta Psychologica, 127*, 449–458.

Robbins, P., Aydede, M. (Eds.). (2009). Cambridge handbook of situated cognition. New York: Cambridge University Press.

Sanford, A. J., & Garrod, S. C. (1981). *Understanding written language: Explanations of comprehension beyond the sentence*. Chichester: Wiley.

Smith, E. R., & Semin, G. R. (2004). Socially situated cognition: Cognition in its social context. *Advances in Experimental Social Psychology, 36*, 53–117.

Smith, W. J. (1977). *The behavior of communicating: An ethological approach*. Cambridge, MA: Harvard University Press.

Stone, J. V. (1999). Object recognition: View-specificity and motion-specificity. *Vision Research, 39*, 4032–4044.

Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, J. E. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science, 268*, 1632–1634.

Tomikawa, S. A., & Dodd, D. H. (1980). Early word meanings: Perceptually or functionally based. *Child Development, 51*, 1103–1109.

Tucker, M., & Ellis, R. (1998). On the relations between seen objects and components of potential actions. *Journal of Experimental Psychology: Human Perception and Performance, 24*, 830–846.

Tucker, M., & Ellis, R. (2004). Action priming by briefly presented objects. *Acta Psychologica, 116*, 185–203.

Vallée-Tourangeau, F., Anthony, S. H., & Austin, N. G. (1998). Strategies for generating multiple instances of common and ad hoc categories. *Memory, 6*, 555–592.

Vuong, Q. C., & Tarr, M. J. (2004). Rotation direction affects object recognition. *Vision Research, 44*, 1717–1730.

Vygotsky, L. S. (1991). Genesis of the higher mental functions. In P. Light, S. Sheldon, & M. Woodhead (Eds.). *Learning to think: Child development in social context* (Vol. 2, pp. 32–41). London: Routledge.

Wimsatt, W. C. (1972). Teleological and logical structure of function statements. *Studies in History and Philosophy of Science, 3*, 1–80.

Wright, L. (1973). Functions. *Philosophical Review, 82*, 139–168.

Yeh, W., & Barsalou, L. W. (2006). The situated nature of concepts. *American Journal of Psychology, 119*, 349–384.

Zwaan, R. A., & Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological Bulletin, 123*, 162–185.