Perceptual Simulation in Conceptual Tasks

Lawrence W. Barsalou
Karen Olseth Solomon
and
Ling-Ling Wu

1. Theoretical framework

In developing a theory of knowledge, a fundamental decision must be made about the nature of representation: Should representations be modal or amodal? We address this important distinction briefly and then summarize several lines of research that bear on it empirically.

1.1. Amodal symbol systems

Across the cognitive sciences, standard theories of knowledge typically adopt the assumptions of what we call *amodal symbol systems.* According to this view, information in the physical world produces neural states in perceptual systems. A transduction process takes these states as input and produces descriptions of them in a completely new representation language. On perceiving a chair, for example, descriptions are transduced that bear no commonalities with their perceptual origins. This lack of commonality exists because the symbols that constitute descriptions are amodal and arbitrary. These symbols are amodal because their structure is unrelated to the structure of perceptual states. These symbols are arbitrary, because they are linked to perceptual states via arbitrary conventions of association. Analogous to how words typically bear arbitrary relations to entities in the world, amodal symbols bear arbitrary relations to the perceptual states that produce them. Just as the phonological structure of a word usually provides no guidance about reference, neither does the structure of an amodal symbol. As a further consequence, the similarity between amodal symbols bears no relation to the similarity between perceptual states (e.g., symbols for *blue* and *green* are not necessarily more similar than symbols for *blue* and *red*).

As amodal symbols become transduced from perceptual states, they enter into larger representational structures, including feature lists, frames, schemata, semantic networks, and logical expressions. In turn, these structures support all of the higher cognitive functions, including knowledge, memory, language, and thought. For reviews of such theories, see Smith and Medin (1981), Rumelhart and Norman (1988), Barsalou (1992), and Barsalou and Hale (1993). For explicit statements about the importance of symbols being amodal and arbitrary, see Newell and Simon (1972), Fodor (1975), Pylyshyn (1984), Haugeland (1985), and Harnad (1990).

1.2. Perceptual symbol systems

A very different approach to knowledge takes the form of what we call *perceptual symbol systems*. According to this view, perceptual states are not transduced into a completely new representational language. Instead, subsets of perceptual states are extracted to function symbolically and support the higher cognitive functions. On perceiving a chair, information is extracted from perceptual representations and transferred to memory. Once in memory, these extractions function symbolically, standing for referents in the world and entering into all forms of symbolic computation. Most importantly, these symbols are modal and analogical, not amodal and arbitrary. These symbols are modal because they have the same structure as perceptual states. These symbols are analogical, first, because their structure is informative about reference (although neither necessary nor sufficient; Goodman, 1976), and, second, because similarity between them corresponds to similarity between perceptual states (assuming a constrained metric of similarity, Goodman, 1955).

This type of theory has been around a long time, at least since Aristotle (4th century BC/1961), and certainly since Locke (1690/1959), Berkeley (1710/1982), Hume (1739/1978), Kant (1787/1965), Reid, (1764/1970, 1785/1969), and, more recently, Russell (1919/1956) and Price (1953). Indeed, until the early twentieth century, nearly all theorists concerned with the mind assumed that knowledge had a strong perceptual character to it. In the modern era, however, most researchers have embraced amodal symbol systems in one form or another. The exceptions have mostly been cognitive linguists, who have made the perceptual character of knowledge a central assumption (e.g., Fauconnier, 1985; Jackendoff, 1987; Johnson, 1987; Lakoff, 1987, 1988; Lakoff & Johnson, 1980; Langacker, 1986, 1987, 1991, Sweetser, 1990; Talmy, 1983, 1988). Outside of linguistics, other cognitive scientists have rediscovered this approach as well. In psychology, these researchers include Mandler (1992), Tomasello (1992), Gibbs (1994), Glenberg (in press), and ourselves (Barsalou, 1993; Barsalou & Prinz, in press; Barsalou, Yeh, Luka, Olseth, Mix, & Wu, 1993; Prinz & Barsalou, in press). In philosophy, these researchers include Barwise and Etchemendy (1990, 1991), Thagard (1992), and Peacocke (1992). In artificial intelligence, see Glasgow (1993). Many additional researchers have considered the role of perceptual simulation in mental imagery, but we focus on its role in knowledge here.

Our work develops perceptual symbol systems in two ways: First, we have constructed a theory of perceptual symbols that supports the standard symbolic functions. Second, we have established an empirical case for this theory. In other papers, we present the theory in greater detail (Barsalou, 1993, 1996; Barsalou & Prinz, in press; Barsalou et al., 1993; Prinz & Barsalou, in press). Although we focus on our empirical work here, we frame it with a brief summary of the theory's seven core assumptions (see Barsalou, 1996, and Barsalou and Prinz, in press, for more complete presentations).

(1) Perceptual symbols are brain states in perceptual systems. We assume that brain states in perceptual systems—not conscious subjective images—constitute perceptual symbols. Although conscious images may accompany the processing of perceptual symbols, much evidence indicates that images are not necessary and that the processing of perceptual symbols often occurs unconsciously. Defining perceptual symbols this way also avoids various problems associated with conscious images, and it provides powerful new tools for representing knowledge perceptually.

(2) Perceptual symbols are schematic. Similar to cognitive linguists (e.g., Langacker, 1986; Talmy, 1983), we assume that perceptual symbols are schematic, containing only small fragments of perceptual states. The basic cognitive mechanism of selective attention extracts a very small subset of the information in a perceptual state and transfers it to long-term memory for later use as a symbol. The resulting symbol is schematic, containing only the extracted information (e.g., the shape of an entity, filtering out its color,

position, etc.).

(3) Perceptual symbols are multimodal. Perceptual symbols do not just represent vision, space, or the sensory modalities. Instead, they can represent any aspect of experience, including all five sensory modalities, proprioception, and introspection. Perceptual symbols extracted from introspection are particularly important (e.g. cognitive states, operations, and affects), because they play central roles in our accounts of abstract concepts. In general, we assume that the schematization process just described extracts perceptual memories from all aspects of experience and stores them for later use as symbols.

(4) Perceptual symbols enter into simulation competence. Perceptual symbols do not reside independently of one another in memory. Instead, they become integrated into systems that allow the cognitive system to simulate entities and events in their absence. On this view, having an adequate concept of something is having an acceptable degree of competence in simulating it. We do not assume that these simulations are ever complete, nor that they are completely veridical or unbiased. Furthermore, a simulation competence is not simply a collection of "sense data." Instead, innate knowledge about space, time, events, objects, and properties structures the perceptual symbols in a simulation competence.

(5) Perceptual symbols are productive. Similar to cognitive linguists (e.g., Langacker, 1986), we assume that a finite number of perceptual symbols can be used to construct an infinite number of complex symbols using combinatoric and recursive mechanisms. Productivity is essentially the inverse process of symbol formation: Whereas information is filtered out of perceptual representations to form schematic perceptual symbols, information is added back to these schematic representations to implement productivity. Because the information added back can go beyond the information filtered, productivity enables the construction of representations for entities and events never experienced.

(6) Perceptual symbols represent propositions. Using the productive system of perceptual symbols developed thus far, it is possible to represent propositions that describe and interpret situations. When attempting to map simulation competencies into a perceived or imagined situation, type-token relations result that can construe a situation in an infinite number of ways that are true or false. This perceptually-based approach implements the basic

functions of propositions in a psychological context.

(7) Perceptual symbols represent abstract concepts directly. Unlike some cognitive linguists, we do not believe that metaphor constitutes the primary representation of abstract concepts such as *truth*, *negation*, *disjunction*, and *anger*. Instead, we believe that most abstract concepts are understood directly in terms of the relevant perceptual experience. Metaphor, when it applies, guides, embellishes, and construes this experience. In our attempts to represent abstract concepts directly, we have found that three mechanisms are central (1) event simulations that frame focal concepts (e.g., Fillmore, 1985; Langacker, 1986), (2) perceptual symbols from introspection, (3) propositional construal. Our conjecture is that any abstract concept can be represented directly using these three mechanisms.

2. Empirical Research

The primary goal of our empirical work has been to test the assumption that perceptual simulation underlies conceptual processing. Whereas standard views assume that people process concepts using amodal symbol systems such as feature lists, frames, and semantic networks, we explore the possibility that people process concepts by simulating their referents perceptually. Although these simulations may at times be conscious, they may often proceed unconsciously (1.2). Note that our empirical efforts thus far do not constitute a full test of our theory. Indeed, most aspects of the theory remain untested. Instead, our research has only explored the theory's core assumption, namely, that perceptual simulation lies at the heart of conceptual processing.

To explore this issue, we have subjects perform standard conceptual tasks, such as feature listing and property verification. In the critical conditions, subjects do not see pictorial stimuli, nor are they asked to perform perceptual processing, such as imagery. Instead, subjects only receive linguistic stimuli and are asked in as neutral a manner as possible to perform conceptual processing. Of interest is whether subjects perform perceptual simulation spontaneously.

2.1. Instructional equivalence and perceptual work

Across several lines of experimentation, we have sought two general forms of evidence for perceptual simulation: *instructional equivalence* and *perceptual work.* To examine instructional equivalence, each experiment includes *neutral* and *imagery* subjects. As just described, neutral subjects receive standard task instructions with nothing said about images. For feature listing, subjects are asked to list the properties typically true of a concept. For property verification, subjects are asked to verify whether a concept has a property. In contrast, imagery subjects are asked to perform these conceptual tasks using images. For feature listing, subjects are asked to construct images for the referents of concepts and then describe these images. For property verification, subjects are asked to construct an image of a concept and then verify the truth or falsity of a property by attempting to find it on the image.

To assess instructional equivalence, we compare detailed performance profiles of the neutral and imagery subjects. We can assume that imagery subjects use images as instructed, given that much previous work indicates that subjects adopt images when asked to do so (for reviews, see Finke, 1989, and Kosslyn, 1980). As discussed shortly, we also have independent verification from perceptual work that imagery subjects use images. Thus, the key issue is how neutral subjects compare to imagery subjects: If neutral subjects use images during perceptual simulation, then their detailed performance profiles should be essentially the same as those for imagery subjects. If both groups perform the task similarly, it is likely that neutral subjects use images, assuming that imagery subjects use them. On the other hand, if neutral subjects use amodal representations such as feature lists, frames, or semantic nets, then the detailed performance profiles of the two groups should differ considerably.[1]

Turning to perceptual work, we manipulate a variable in each experiment that is known to affect perceptual processing, such as the visibility, size, or position of features. In both perception and imagery, it is well known that these variables affect performance (Finke, 1989; Kosslyn, 1980; Shepard & Cooper, 1982). If neutral subjects in our conceptual tasks perform perceptual simulation, then we should see effects of these variables on their performance. The more perceptual work a subject has to do, the more difficult processing should be. In contrast, the amodal view predicts no such effects because perceptual simulation is irrelevant.

Perceptual work provides an independent check on whether imagery subjects use images as instructed. If they do, then perceptual work variables should affect their performance. To the extent that perceptual work variables affect neutral and imagery subjects similarly, instructional equivalence is established further.

2.2. Feature listing

In her dissertation, Ling-Ling Wu explored instructional equivalence and perceptual work in the feature listing task (Wu, 1995). In these experiments, subjects received nouns or noun phrases and were asked to produce features for them verbally. Wu performed detailed analyses of subjects' tape-recorded protocols to assess the hypotheses of interest. Whereas all previous accounts of feature listing assume that subjects produce features from feature lists, frames, or semantic networks, Wu predicted that subjects would simulate referents of these concepts perceptually and scan across them to identify features.

2.2.1. Method and predictions. To assess instructional equivalence and perceptual work, Wu manipulated two orthogonal variables between subjects. For instructional equivalence, Wu manipulated production mode by giving subjects imagery, neutral, or word association instructions. As described earlier, the imagery subjects were asked to "construct an image for the concept and describe it", whereas the neutral subjects were simply asked to "list the characteristics typically true of a concept." The word association subjects were asked to "produce associated words that come to mind for a concept."

According to the perceptual view, neutral and imagery subjects should produce the same detailed distributions of features (instructional equivalence), because both groups use perceptual simulation to generate features. Furthermore, neutral and imagery subjects should differ from word association subjects, who search lexical networks to produce features instead. In contrast, the amodal view predicts that neutral and imagery subjects should differ, because neutral subjects access features from feature lists, frames, and/or semantic nets, whereas imagery subjects access features via perceptual simulation. Furthermore, neutral subjects should perform similarly to word association subjects, because the amodal structures that neutral subjects access are often presumed similar to the linguistic structures that word association subjects access.

For perceptual work, Wu manipulated concept complexity. Half of the eight subjects in each production group received ten critical nouns, and the other half received the same nouns with revealing modifiers. For both groups, the nouns specified concepts such as *watermelon*, *face*, and *lawn*. In the noun phrase group, revealing modifiers were added to form complex concepts such as *half watermelon*, *smiling face*, and *rolled-up lawn*. When unmodified, these nouns refer to entities whose internal features are occluded by external surfaces (e.g., a green rind occludes *red* and *seeds* for *watermelon*). In contrast, when these nouns are coupled with revealing modifiers, the referents are entities whose internal parts are revealed (e.g., *red* and *seeds* for *half watermelon*).

According to the perceptual view, neutral and imagery subjects should produce more internal features for the noun phrases than for the nouns, because perceptual simulations for the noun phrases reveal the insides of entities. Conversely, neutral and imagery subjects should produce few internal features for the nouns, because these perceptual simulations occlude internal features. The amodal view, as typically construed, makes no such prediction. Because amodal accounts of conceptual combination typically assume compositionality, a modifier should only affect the attribute it modifies and no other part of the concept. For *half watermelon, half* should only affect *amount* and have no bearing on internal features. Thus, compositional versions of the amodal view predict no effect of concept complexity. Non-compositional versions can be developed, but they are post hoc and essentially approximate the perceptual view, as we shall see later.

2.2.2. Coding analysis. To analyze the protocol that a subject produced for a concept, Wu first parsed it into elemental features. She then coded each feature into one of 34 types and each cluster of related features into one of 9 types. High levels of coding reliability were established for parsing, features, and clusters (always above .8 and usually above .9).

The 34 coding categories for individual features were organized hierarchically into 4 general types: entity, situation, introspection, and taxonomy. Within each general type, more specific types were defined, which we illustrate with examples for *watermelon*. Entity features included external components (*rind*), external perceived properties (*green*), internal components (*seeds*), internal perceived properties (*sweet*), and six additional types. Situation features included locations (*park*), participants (*family*), activities (*eating*), and nine additional types. Introspective features included affects/emotions (*like it*),

Table 1. Correlations between neutral, imagery, and word association subjects across feature and cluster frequencies (corrected correlations in parentheses).

	Neutral-Imagery Correlations		Neutral-Word Correlations	
	Nouns	Noun phrases	Nouns	Noun phrases
Features	.89 (.99)	.96 (.99)	.73 (.89)	.29 (.39)
Clusters	.81 (.93)	.94 (.99)	.61 (.63)	.15 (.18)

cognitive operations (*similar to cantaloupe*), negations (*not neat and tidy*), and four additional types. Taxonomic features included ontological categories (*object*), superordinates (*fruit*), coordinates (*cantaloupe*), and three additional types. Although subjects were only instructed to produce entity features (e.g., "list characteristics typically true of watermelons"), they also produced extensive numbers of situation, introspective, and taxonomic features, indicating that they situated objects environmentally, introspectively, and taxonomically.

The nine coding categories for feature clusters were assigned to two or more contiguous features that described a common entity, the external surface of an entity, the internal region of an entity, a specific situation, a generic situation, an introspection, a time period, or an attribute's values. These eight cluster types could embed within one another hierarchically. The ninth cluster type was for one or more features that could not be meaningfully related to surrounding features.

2.2.3. Results. To assess instructional equivalence, Wu computed the average frequency of each feature type per protocol in each of the six between-subjects conditions. She similarly computed the average frequency of each cluster type. In a given condition, the averages for the 34 feature types constitute a complex performance profile for the types of information that subjects produced. Similarly, the averages for the 9 cluster types constitute a complex performance profile for how subjects organized their protocols. These two sets of performance profiles were used to assess instructional equivalence. If neutral subjects construct and scan images to produce features, their performance profiles should correlate highly with those for imagery subjects.

Table 1 presents the results from this analysis. For neutral and imagery subjects, the correlations across frequencies for the 34 feature types were very high. In contrast, the correlations between neutral and word association subjects

Table 2. Average frequency of internal features and internal clusters in nouns and noun phrases (NPs) (average positions in parentheses).

	Neutral		Imagery		Words	
	Nouns	NPs	Nouns	NPs	Nouns	NPs
Features	.90 (.54)	2.23 (.42)	1.15 (.68)	2.45 (.44)	.30 (.65)	.70 (.52)
Clusters	.28 (.53)	.68 (.40)	.23 (.70)	.80 (.50)	.03 (.50)	.18 (.57)

were much lower. The results for feature clusters were similar. Furthermore, when the correlations between imagery and neutral subjects were corrected for reliability, they all rose well above .90, with three being .99 (see Table 2). These results indicate that neutral and imagery subjects produced similar features and feature organizations. Thus, instructional equivalence occurred to a considerable extent, suggesting that neutral subjects constructed and scanned images to produce features. The fact that neutral and word association subjects produced substantially different protocols indicates, first, that the production manipulation has effects, and second, that neutral subjects did not use representations that were very language-like.

Turning to perceptual work, if neutral and imagery subjects simulated concepts perceptually, they should have produced more internal features for the noun phrases than for the nouns, because internal features should be revealed for the noun phrases but occluded for the nouns. The results in Table 2 support this prediction strongly. Subjects produced significantly more internal features for noun phrases than for nouns, not only in the imagery condition, but also in the neutral condition. Subjects also produced significantly more internal clusters for noun phrases than for nouns in both conditions. Consistent with the perceptual view, the increased frequency of internal features and internal clusters for noun phrases demonstrates an increased salience of internal information.

Also consistent with the perceptual view, internal features occurred earlier in the protocols for noun phrases than for nouns. Using a normalized position measure, where 0 indicates the first feature in a protocol and 1 indicates the last, subjects produced internal features earlier for noun phrases than for nouns in both the neutral condition and the imagery condition (see Table 2). Subjects also produced internal clusters earlier for noun phrases than for nouns in both conditions. The earlier positions of internal features and internal clusters in noun phrases further demonstrates the increased salience of internal information.

Two additional results further implicate perceptual simulation: First, most subjects reported having images to a substantial extent, even in the word association conditions. Second, subjects produced surface features earlier than external components. This is consistent with an implication of the perceptual view that the more external a feature is, the earlier it should be reported. Because surface features lie only on the surface (e.g., *green* for *watermelon*), whereas surface components extend below it (e.g., *rind*), subjects should report surface features earlier, and they do for both nouns and noun phrases.

2.2.4. Further experiments. In her dissertation, Wu performed two additional experiments that replicate and extend these results. In Experiment 2, she observed instructional equivalence and perceptual work, not only for familiar noun phrases (e.g., *half watermelon, open computer, dug-up lawn*), but also for novel noun phrases (e.g., *gashed watermelon, transparent computer, rolled-up lawn*). In Experiment 3, she observed instructional equivalence and perceptual work for both externally-focused noun phrases (e.g., *striped watermelon, shiny car, scratched dresser*) and internally-focused noun phrases (e.g., *seedless watermelon, comfortable car, disorganized dresser*). For the externally-focused modifiers, external features became more salient than internal features; for the internally-focused modifiers, internal features became more salient. These results confirm a priori predictions of the perceptual view that features salient in a perceptual simulation should be reported earlier and more often.

In a final experiment, to be included in a published report of this work, we found that revealing modifiers are noun dependent. In the revealing condition, neutral subjects produced features for concepts such as *rolled-up lawn*, *comfortable car*, and *open computer*, with the prediction again being that internal features should be salient, because they are revealed in referents. In contrast, neutral subjects in the non-revealing condition produced features for concepts such as *rolled-up snake*, *comfortable watch*, and *open hands*, with the prediction being that internal features should not be salient, because they are not revealed in simulations of these referents. The results supported this prediction, indicating that amodal rules for modifiers can't account for the results of the previous experiments (e.g., "if *rolled-up* modifies a noun, increase the salience of internal features in the noun's concept"). Instead, an amodal view requires a different post hoc rule for every noun concept, which rests on a perceptual examination of its referents. A more parsimonious and principled account is that subjects simulate referents perceptually, with features being more salient when revealed than occluded.

2.3. Property verification

Perhaps the perceptual processing in Wu's experiments reflects the fact that feature listing is a slow, sequential, and recall-oriented task. Because subjects have no time constraints, they can generate detailed images and use them to produce features. If so, then we might not see perceptual processing in a faster, recognition-oriented task. If perceptual simulation underlies conceptual processing in general, however, we should observe it in conceptual tasks of both types. To explore this issue, Karen Olseth Solomon has been assessing instructional equivalence and perceptual work in the property verification task (for a preliminary report of this work, see Olseth & Barsalou, 1995).

In the property verification task, subjects receive the name of a concept followed by the name of a property and have to verify whether the property is true of the concept. For example, subjects might verify that *mane* is a property of *horse*, or they might verify that *faucet* is not. Subjects make few errors (under 10%) and generally produce reaction times under 1 sec.

Theoretical accounts assume that subjects verify properties by retrieving concept-property relations from feature lists, frames, semantic nets, and so forth. The only dissenting voice to the amodal view has been Kosslyn (1976, 1980). When Kosslyn asked some subjects to find the properties on images of the concepts, he found that property size predicted verification time: The larger the property, the faster the verification. Based on this result, Kosslyn concluded that subjects can use perceptual representations in this conceptual task, with larger properties being easier to find on images. However, Kosslyn also ran a group of neutral subjects, whose performance was not influenced by property size. Because associative strength influenced their performance instead, Kosslyn concluded that neutral subjects use amodal semantic nets, with highly associated properties being accessed faster from a concept's node than weakly associated properties. Because these properties are represented amodally, size has no effect on verification time.

A potential problem with this study is that neutral subjects may not have performed conceptual processing, thereby precluding any conclusions about the nature of conceptual representations. To see why this could be a problem, consider the false trials of the experiment. On these trials, the properties appeared quite easy to reject, because they were not associated with their respective concepts (e.g., *MOUSE-stinger*). In contrast, the true properties tended to be associated with their concepts (e.g., *CAT-claw*). Thus, subjects could have adopted the following strategy: If the words for the concept and property are associated, respond true; if they're unassociated, respond false. Subjects could have simply used linguistic associations between words to perform the task, rather than accessing conceptual knowledge about objects and their properties. If subjects adopted this associativeness strategy, it precludes a test of whether conceptual knowledge is amodal or perceptual.

To remedy this problem, Solomon performed an experiment in which she manipulated the associativeness of the false properties. Similar to Kosslyn's experiments, the false properties were unassociated with their concepts for half of the subjects (e.g., CRAB-brick, BEAVER-rope, MONKEY-string). However, the false properties for the remaining subjects were highly associated with their concepts (e.g., CRAB-fin, BEAVER-quills, MONKEY-banana).

When false properties are unassociated, Solomon predicted lack of instructional equivalence, similar to Kosslyn. Because neutral subjects can use the associativeness strategy, they should not access conceptual knowledge and should therefore not exhibit perceptual work. In contrast, imagery subjects should exhibit perceptual work, because they follow instructions and use perceptual representations. When false properties are associated, however, Solomon predicted both instructional equivalence and perceptual work. Because these false properties block the associativeness strategy, neutral subjects must access conceptual knowledge. If this knowledge is perceptual, then these subjects, like imagery subjects, should exhibit perceptual work.

2.3.1. Method. To assess these predictions, Solomon manipulated two variables orthogonally between subjects: instructions and associativeness. Analogous to Kosslyn's studies, half of the subjects received imagery instructions, and half received neutral instructions. As just described, half of the subjects in each instructional condition received unassociated false properties, and half received associated false properties. Independent scaling verified the associativeness manipulation.

Following 76 practice trials, 44 subjects performed 224 critical trials, half true and half false in a random order. Concepts ranged widely across natural kinds and artifacts. Each property was only seen once, although concepts were presented on multiple trials (on the average, about twice on true trials and twice on false trials). To begin each trial, subjects pressed a foot pedal and received the word for a concept for 500 msec. After a blank screen for 1200 msec, the word for a property appeared simultaneously with the initiation of the reaction time measurement in milliseconds, which terminated with a true or false key press. Subjects were instructed to respond true only if the property was a physical part of the concept's referents. Feedback was not provided about correctness of the response. Errors were less than 10%, there was no speed-accuracy trade-off, and unusually long reaction times were removed from the analysis.

2.3.2. Scaling and regression. To test the hypotheses, it was necessary to scale the true concept-property pairs on a wide variety of measures that fell into four groups: associative, perceptual, property expectancy, and property goodness. The associative measures included concept-to-property strength, property-to-concept strength, and two others. The perceptual measures included the absolute area of properties in the image field, the absolute length of properties, the position of properties relative to the top of the image, the position of properties relative to the left of the image, the salience of properties in the image, and seven others. The property expectancy measures included how many objects have a similar property, how similar the property is across different objects, and one other. The property goodness measures included whether the properties were components, regions, or materials, and whether the word for the property is more likely to be interpreted as an object than as a property (e.g., *barrel* for *rifle*). The data for all measures were collected from independent groups of subjects, none who participated in the main experiment. Solomon scaled every measure that struck us as having any potential for capturing variance in the verification task. Besides the twenty measures in the final analyses, Solomon scaled another twenty or so that turned out to be redundant with other measures. Once the scaling process was completed, Solomon had a value on each scaled measure for each true concept-property pair. Because the associative and perceptual measures were of primary theoretical interest, and because they accounted for most of the variance, we focus on them here, excluding the property expectancy and property goodness measures.

To test the key hypotheses, Solomon regressed the mean reaction times for the true concept-property pairs onto the associative and perceptual measures, once for each of the four instruction by associativeness conditions. Of interest was which measures best predicted verification times in each condition. Most importantly for the perceptual view, perceptual factors should predict performance in the neutral-associated condition. Because the associated false properties block the associativeness strategy, neutral subjects should perform perceptual simulation to verify properties, such that perceptual factors predict verification times. In contrast, when the false properties are unassociated, neutral subjects can adopt the associativeness strategy, such that associative factors dominate prediction.

To ensure broad conclusions, Solomon performed these regressions in many manners, with the qualitative pattern of results remaining constant across them. The regressions reported here are hierarchical, with the associative measures always entered into the regressions before the perceptual measures. Thus, the perceptual measures only account for additional variance that the associative measures did not capture first. This provides a conservative test of whether perceptual factors play a central role in property verification.

2.3.3. Results and discussion. Table 3 presents the results from this analysis. When the false properties were unassociated with their concepts, associative measures explained the most variance for neutral subjects, and perceptual measures did not account for significant variance. As predicted, these subjects adopted the associativeness strategy, because the associative strength between concepts and properties provided useful information for producing correct responses.

In contrast, when the false properties were associated with their concepts, associative measures accounted for no significant variance in the neutral condition. Because both true and false properties were associated with their concepts, subjects could not use associativeness as a cue but had to use conceptual knowledge. Most importantly, this conceptual knowledge appears to be perceptual, because perceptual measures now account for the most variance. Even though the perceptual measures entered the regression after the associative measures, they accounted, not only for significant unique variance, but for the most. Thus, when neutral subjects are forced to perform conceptual processing, they shift from an associative strategy to perceptual simulation.

Table 3. Percentage of variance accounted for in a hierarchical regression, with associative factors entered first.

	Unassociated false trials		Associated false trials	
Factors	Neutral	Imagery	Neutral	Imagery
Associative	16*	16*	6	6
Perceptual	10	15*	18*	29*

* Significant at p < .05 in accounting for unique variance.

The two imagery conditions further support these conclusions. In the imagery-unassociated condition, both associative and perceptual measures accounted for significant variance. Because these subjects followed instructions and used images, perceptual measures predicted their performance. Because the unassociated false properties allowed the associativeness strategy, associative strength predicted performance as well. In contrast, associative strength did not account for significant variance in the imagery-associated condition, indicating that the associated false properties again blocked the associativeness strategy. Instead, only the perceptual measures accounted for significant variance.

In summary, this experiment found evidence of both instructional equivalence and perceptual work. When associated false trials blocked the associativeness strategy, instructional equivalence held between neutral and imagery subjects. Furthermore, the dominant roles of perceptual measures in both conditions implicate perceptual work. Thus, instructional equivalence and perceptual work don't just occur for feature listing—they also occur for property verification. Perceptual simulation is central in conceptual tasks that are fast and recognition-oriented, as well as in ones that are slow and recall-oriented.

One other result deserves mention. As described earlier, Kosslyn (1976, 1980) found that imagery subjects verified large properties faster than small ones, suggesting that it's easier to find large properties on images. In contrast, we found that small properties were verified faster than large properties in every condition, holding all other factors constant. We interpret this result as indicating that it takes longer to construct a large property on an image than a small one, consistent with Kosslyn, Reiser, Farah, and Fliegel (1983).

2.3.4. Further experiments. As Figure 3 illustrates, the total variance accounted for is not high. This primarily reflects the low reliability of the average reaction times for the concept-property pairs. A later experiment in Solomon's dissertation attempts to replicate these results, using more subjects to increase reliability and explained variance. This experiment also uses materials that do not repeat concepts, and it uses new scaling methods that appear superior to the original ones.

2.4. Part similarity

The finding that subjects simulate a part of an object in order to verify it leads to a further prediction: If the same part is constructed on images of two different objects, the ease of constructing the part on the second object should depend on the nature of the part constructed on the first object. To see this prediction, imagine that a subject verifies *mane* for *pony* at one point in an experiment. Imagine on a later trial that the subject verifies *mane* for *horse* and that the earlier image of a pony's mane becomes active. Because the image of the pony's mane is much like the mane that needs to be constructed on the image of a horse, the mane for *pony* facilitates this process. In contrast, imagine that a different subject first verifies *mane* for *lion* and later verifies *mane* for *horse*. The mane of a lion is qualitatively different from the manes of ponies and horses. Whereas a lion's mane traverses the circumference of its neck, the mane of a pony or horse traverses the length of its neck. Thus, verifying *mane* for *lion* could have a detrimental effect on later verifying *mane* for *horse*.

To test this hypothesis, Solomon and Barsalou (1996) had subjects verify target pairs (e.g., *HORSE-mane*) either after verifying similar parts earlier in the experiment (e.g., *PONY-mane*) or after verifying dissimilar parts (e.g., *LION-mane*). The perceptual view predicts that the similarity of a part's two instantiations will affect performance on the second verification. In contrast, standard amodal views do not make this prediction a priori. Because they represent *mane* using the same part concept in all object concepts, they predict no difference. Amodal views could be developed to account for this result, but they do so post hoc and require perceptual representations to make the prediction.

A potential confounding is that concepts in the similar-parts condition are more similar taxonomically than in the different-parts condition. For example, *pony* and *horse* are more similar than *lion* and *horse*. Thus, faster processing for

HORSE-mane after *PONY-mane* than after *LION-mane* could reflect the taxonomic similarity of the concepts, not the perceptual similarity of the parts. To control this confounding, a second property was found that was approximately the same across all three concepts. For example, *back* is approximately the same for *horse*, *pony*, and *lion*. If taxonomic similarity is important, subjects should verify *HORSE-back* faster after verifying *PONY-back* than after verifying *LION-back*. Conversely, if part similarity is important, there should be no difference, because *back* is equally similar for all three concepts. Independent scaling verified this manipulation.

2.4.1. Method. Two orthogonal factors were manipulated within subjects: Concept similarity and part similarity. Concept similarity concerned the similarity of the two concepts in a sequence, either high (e.g., *PONY* followed by *HORSE*) or low (e.g., *LION* followed by *HORSE*). Part similarity concerned whether a property was similar across all three concepts (e.g., *back* for *HORSE*, *PONY*, and *LION*), or similar for only two of them (e.g., *mane* for *HORSE* and *PONY* but not *LION*). There were 16 total sets of such materials. Each of the 48 subjects received only one sequence of pairs from each set, but received four sequences of each type across sets in a carefully controlled and counterbalanced manner. The number of trials intervening between the two pairs of a sequence ranged from 10 to 25 with an average of 17.5. Filler trials obscured the structure of the critical materials. The procedure for timing the trial sequence was the same as in the previous experiment (2.3.1), except that subjects received feedback.

2.4.2. Results and discussion. The results in Table 4 confirm the predictions of the perceptual view. For parts that differed across the three concepts (e.g., *mane*), subjects were significantly faster to verify the second pair if the earlier instantiation of the part was similar than if it was dissimilar. This effect did not reflect the taxonomic similarity of the concepts. For parts that were the same across all three concepts (e.g., *back*), subjects were no faster to verify sequences of similar concepts (e.g., *PONY* followed by *HORSE*) than sequences of dissimilar concepts (e.g., *LION* followed by *HORSE*).

Table 4. Average reaction time in milliseconds for the second concept-property pair of a sequence (error rates in parentheses).

Condition	Example of a sequence		RT (errors)
Similar parts	PONY-mane	HORSE-mane	725 (2%)
Dissimilar parts	LION-mane	HORSE-mane	778 (5%)
Similar categories	PONY-back	HORSE-back	691 (4%)
Dissimilar categories	LION-back	HORSE-back	695 (2%)

Thus, the critical factor was the perceptual similarity of the parts. Processing a part's second instantiation was faster when its first instantiation was similar than when it was dissimilar. When part similarity was held constant, taxonomic similarity had no effect, indicating that only part similarity was important. These results confirm a strong a priori prediction of the perceptual view that the perceptual detail of parts can enter into conceptual processing.

2.4.3. Further experiments. Because the first experiment had an interval of 1700 msec between the onsets of the concepts and the properties, various strategic processes could have entered into subjects' performance. To eliminate this possibility, a second experiment shortened the onset asynchrony between the concept and property to 250 msec. As has been well established in the literature, this short interval precludes any strategic processing that could bias performance (e.g., Neely, 1977). This change in the procedure had no effect, with the same pattern emerging as in the first experiment. In two other experiments, we are implementing a variety of base line conditions to establish whether the critical effects reflect facilitation and/or interference.

3. Conclusions

In this paper, we have tried to make two fundamental points: First, it is possible to develop a powerful and sophisticated theory of perceptual symbols that supports the standard symbolic functions. Second, strong empirical evidence exists that perceptual simulation underlies conceptual processing.

Acknowledgment and Footnote

The research reported in this chapter was supported by National Science Foundation Grant SBR-9421326 to the first author. Address correspondence to Lawrence W. Barsalou, Department of Psychology, 5848 S. University Ave., University of Chicago, Chicago, IL 60637, L-Barsalou@uchicago.edu.

[1] The reader might worry that we predict the null hypothesis, namely, that the neutral and imagery groups should not differ. However, we always include either a third group of subjects, whom we predict to differ from neutral and imagery subjects, or a variable on which imagery and neutral subjects should differ on one level. Thus, we always predict complex patterns of results that include positive effects. Furthermore, the perceptual and amodal views make complex patterns of predictions across the manipulations of instructional equivalence and perceptual work together. The perceptual view predicts a null effect for instructional equivalence but a positive effect for perceptual work, whereas the amodal view predicts a positive effect for instructional equivalence but a null effect for perceptual work. Again, each view makes complex patterns of prediction that go considerably beyond local predictions of null effects.

References

Aristotle
1961 *De anima, books II and III* (Hamlyn, D. W., Trans.). Oxford: Oxford University Press. (Original work published in 4th century B.C.)

Barsalou, L.W.
1992 Frames, concepts, and conceptual fields. In E. Kittay & A. Lehrer (Eds.), *Frames, fields, and contrasts: New essays in semantic and lexical organization* (21-74). Hillsdale, NJ: Erlbaum.

Barsalou, L.W.
1993 Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In A.C. Collins, S.E. Gathercole, & M.A. Conway (Eds.), *Theories of memories* (29-101). London: Erlbaum.

Barsalou, L.W.
1996 Perceptual symbol systems. Under review.

Barsalou, L.W., & Hale, C.R.
1993 Components of conceptual representation: From feature lists to recursive frames. In I. Van Mechelen, J. Hampton, R. Michalski, & P. Theuns (Eds.), *Categories and concepts: Theoretical views and inductive data analysis* (97-144). San Diego, CA: Academic Press.

Barsalou, L.W., & Prinz, J.J.
in press Mundane creativity in perceptual symbol systems. In T.B. Ward, S.M. Smith, & J. Vaid (Eds.), *Conceptual structures and processes: Emergence, discovery, and change*. Washington, DC: American Psychological Association.

Barsalou, L.W., Yeh, W., Luka, B.J., Olseth, K.L., Mix, K.S., & Wu, L.
1993 Concepts and meaning. In K. Beals, G. Cooke, D. Kathman, K.E. McCullough, S. Kita, & D. Testen (Eds.), *Chicago Linguistics Society 29: Papers from the parasession on conceptual representations.* University of Chicago: Chicago Linguistics Society.

Barwise, J., & Etchemendy, J.
1990 Information, infons, and inference. In R. Cooper, K. Mukai, and J. Perry (Eds.), *Situation theory and its applications* (pp. 33-78). Chicago: University of Chicago Press.

Barwise, J., & Etchemendy, J.

1991 Visual information and valid reasoning. In. W. Zimmerman & S. Cunningham (Eds.), *Visualization in mathematics* (pp. 9-24). Washington: Mathematical Association of America.

Berkeley, G.
1982 *A treatise concerning the principles of human knowledge*. Indianapolis, IN: Hackett. (Original work published in 1710)

Fauconnier, G.
1985 *Mental spaces.* Cambridge, MA: MIT Press.

Fillmore, C.J.
1985 Frames and the semantics of understanding. *Quaderni di Semantica, 6*, 222-255.

Finke, R.A.
1989 *Principles of mental imagery*. Cambridge, MA: MIT Press.

Fodor, J. A.
1975 *The language of thought*. Cambridge, MA: Harvard University Press.

Gibbs, R.W., Jr.
1994 *Poetics of mind.* New York: Cambridge University Press.

Glasgow, J.I.
1993 The imagery debate revisited: A computational perspective. *Computational Intelligence, 9,* 309-333.

Glenberg, A.
in press What memory is for. *Behavioral and Brain Sciences*.

Goodman, N.
1955 *Fact, fiction, and forecast*. Cambridge, MA: Harvard University Press.

Goodman, N.
1976 *Languages of art.* Indianapolis: Hackett.

Harnad, S.
1990 The symbol grounding problem. *Physica D, 42,* 335-346.

Haugeland, J.
1985 *Artificial intelligence: The very idea of it.* Cambridge, MA: MIT Press.

Hume, D.
1978 *A treatise on human nature.* (2nd ed.). Oxford: Oxford University Press. (Original work published in 1739)

Jackendoff, R.
1987 On beyond zebra: The relation of linguistic and visual information. *Cognition, 26,* 89-114.
Johnson, M.
1987 *The body in the mind: The bodily basis of reason and imagination*. Chicago: University of Chicago Press.
Kant, E.
1965 *The critique of pure reason* (N.K. Smith, Trans). New York: St. Martin's Press. [Original work published in 1787]
Kosslyn, S.M.
1976 Can imagery be distinguished from other forms of conceptual representation? Evidence from studies of information retrieval times. *Memory & Cognition, 4, 291-297*.
Kosslyn, S.M.
1980 *Image and mind.* Cambridge, MA: Harvard University Press.
Kosslyn, S.M., Reiser, B.J., Farah, M.J., & Fliegel, S.J.
1983 Generating visual images: Units and relations. *Journal of Experimental Psychology: General*, *112*, .278-303
Lakoff, G.
1987 *Women, fire, and dangerous things: What categories reveal about the mind.* Chicago: University of Chicago Press.
Lakoff, G.
1988 Cognitive semantics. In U. Eco, M. Santambrogio, & P. Violi (Eds.), *Meaning and mental representations* (pp. 119-154).Bloomington, IN: Indiana University Press.
Lakoff, G., & Johnson, M.
1980 *Metaphors we live by*. Chicago: University of Chicago Press.
Langacker, R.W.
1986 An introduction to cognitive grammar. *Cognitive Science*, *10*, 1-40.
Langacker, R.W.
1987 *Foundations of cognitive grammar: Vol. 1. Theoretical prerequisites.* Stanford, CA: Stanford University Press.
Langacker, R.W.
1991 *Foundations of cognitive grammar*, *volume II: Descriptive application.* Stanford, CA: Stanford University Press.
Locke, J.
1959 *An essay concerning human understanding.* (1st ed.). (Vol. I and II). New York: Dover. (Original work published in 1690)
Mandler, J.M.
1992 How to build a baby: II. Conceptual primitives. *Psychological Review, 99,* 587-604.
Neely, J.H.
1977 Semantic priming and retrieval from lexical memory. Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General*, *106*, 226-254.
Newell, A., & Simon, H.A.
1972 *Human problem solving.* Englewood Cliffs, NJ: Prentice-Hall.

Olseth, K.L., & Barsalou, L.W.
1995 The spontaneous use of perceptual representations during conceptual processing. *Proceedings of the seventeenth annual meeting of the Cognitive Science Society* (310-315). Hillsdale, NJ: Lawrence Erlbaum Associates.
Peacocke, C.
1992 *A study of concepts*. Cambridge, MA: MIT Press.
Price, H. H.
1953 *Thinking and experience*. London: Hutchinson's Universal Library.
Prinz, J.J., & Barsalou, L.W.
in press Acquisition and productivity in perceptual symbol systems: An account of mundane creativity. In T.H. Dartnall (Ed.), *Creativity,computation, and cognition.* Cambridge, MA: MIT/AAAI Press.
Pylyshyn, Z.W.
1984 *Computation and cognition*. Cambridge, MA: MIT Press.
Reid, T.
1969 *Essays on the intellectual powers of man*. Cambridge, MA: MIT Press. (Original work published in 1785)
Reid, T.
1970 *An inquiry into the human mind on the principles of common sense*. Chicago, IL: University of Chicago Press. (Original work published in 1764)
Rumelhart, D.E., & Norman, D.A.
1988 Representation in memory. In R.C. Atkinson, R.J. Herrnstein, G. Lindzey, & R.D. Luce (Eds.), *Stevens' handbook of experimental psychology: Vol. 2. Learning and cognition* (pp. 511-587). New York: Wiley.
Russell, B.
1956 *Logic and knowledge*. London: Routledge. (Original work published in 1919)
Shepard, R.N., & Cooper, L.A.
1982 *Mental images and their transformations*. New York: Cambridge University Press.
Smith, E.E., & Medin, D.L.
1981 *Categories and concepts*. Cambridge, MA: Harvard University Press.
Solomon, K.O., & Barsalou, L.W.
1996 Work in progress. University of Chicago.
Sweetser, E.
1990 *From etymology to pragmatics: Metaphorical and cultural aspects of semantics*. Cambridge, England: Cambridge University Press.
Talmy, L.
1983 How language structures space. In H. Pick & L. Acredelo (Eds.), *Spatial orientation: Theory, research, and application* (225-282). New York: Plenum Press.
Talmy, L.
1988 Force dynamics in language and cognition. *Cognitive Science, 12*, 49-100.
Thagard, P.
1992 *Conceptual revolutions*. Princeton, NJ: Princeton University Press.
Tomasello, M.
1992 First verbs: A case study of early grammatical development. New York: Cambridge University Press.
Wu, L.L.
1995 *Perceptual representation in conceptual combination*. Doctoral dissertation. University of Chicago.