

106. Conceptual knowledge, categorization, and meaning

1. Introduction
2. Conceptual knowledge
3. Theoretical approaches
4. Issues of debate
5. Conceptual combination
6. Relationship between conceptual knowledge and word meanings
7. Concluding remarks
8. References

Abstract

Since Eleanor Rosch's groundbreaking work in the 1970s, conceptual knowledge has become a subject of extensive research in cognitive psychology. This chapter provides an overview of the current state of the art. Research has focused on conceptual knowledge about concrete physical things. The main research questions concern the structure and content of conceptual knowledge and its functions, in particular categorization. Most research is based on the view that conceptual knowledge comprises a set of relatively fixed packets of information, or concepts, which are assumed to correspond to lexical meanings. This view of the relationship between conceptual and lexical-semantic knowledge is discussed towards the end of the chapter.

1. Introduction.

The human mind does not have direct access to the world. What is taken as a real situation in the world is the content of a mental representation constructed from sensory data and knowledge stored in long-term memory. Conceptual knowledge plays a pivotal role here, imposing a particular structure on the representation and promoting a conceptualization in terms of entities of particular kinds, possessing certain properties and being related to each other in particular ways. For example, a given dynamic visual input may, by virtue of conceptual knowledge, give rise to the perception of a structured motion event such as a rabbit jumping into the room. In a similar way, conceptual knowledge also shapes the structure and contents of mental representations in thinking and action planning. It is important to distinguish between conceptual knowledge itself and mental representations constructed at certain points in time that are shaped by conceptual knowledge. The distinction may best be framed in terms of the distinction between long-term memory, which is a permanent store of information, and working memory, where temporary representations are created and manipulated. In working memory, information retrieved from long-term memory is integrated with information from the sensory-motor and emotional systems, and the resulting representations are also heavily influenced by motivational factors and attentional processes. Thus, the mental representation of a situation, event, or individual entity currently perceived or thought of is a representation in working memory. It is shaped by conceptual knowledge but is not part of conceptual knowledge. Conceptual knowledge itself is a component of long-term memory.

Language is a means by which a person can convey information residing in working memory to another person. In doing so, the person needs to carve up the working memory representation and package the intended information in a way which conforms to the linguistic structures of his or her language. For example, to communicate the above mentioned motion event of a rabbit jumping into the room, the various pieces of information, including the entities and the manner and path of motion, must be organized in a particular way. Obviously, the difficulty of the task largely depends on how similar the required structure is to the structure of the given working memory representation, as induced by conceptual knowledge. Many cognitive psychologists assume that conceptual knowledge comprises distinct concepts, each of which corresponds to the meaning of a particular lexical item. If so, carving up working memory representations for the purpose of coding their contents linguistically would be a relatively straightforward process. However, matters are far from settled. As yet few studies have addressed the conceptualization of complex situations or actions and their mapping onto linguistic structures, except in research on the linguistic relativity hypothesis (see Sec. 6) and language development (see, e.g., Snedeker & Gleitman 2004; see also article 107 (Landau) *Space in semantics and cognition*). By far the most studies of conceptual knowledge are concerned with concepts of everyday physical things (for research on other noun concepts, see, e.g., Wisniewski 2009 and Papafragou 2005 on substance concepts, and Goldwater, Markman & Stilwell 2011 on relational concepts). Moreover, even for concepts of everyday physical things, the claim that they correspond to lexical meanings is difficult to evaluate. The reason is that research on conceptual knowledge in general simply presupposes that concepts are word meanings, rather than investigating this issue experimentally. Due to this presupposition, it is common practice in empirical studies to employ verbal stimuli to investigate concepts and when using nonverbal tasks, rarely is any effort made to control for internal linguistic processes such as covert naming or priming from preceding or expected linguistic tasks. Thus, for many of the studies on conceptual knowledge it is strictly speaking impossible to decide whether the results do in fact reveal something about conceptual knowledge or rather about lexical semantic knowledge. However, notwithstanding this unfortunate ambiguity, the findings are in any case of interest to semantics.

This chapter provides an overview of research on conceptual knowledge in cognitive psychology. Its focus is on behavioral research (cf. Martin & Caramazza 2003 for neuroscientific research). In Sections 2 to 5 we report empirical findings on the content and structure of conceptual knowledge and outline the different theoretical approaches as well as their major points of contention. In these sections, we adopt the view of concepts as lexical meanings, but in Section 6, we explicitly address the question of how conceptual and lexical-semantic knowledge are related.

The literature on conceptual knowledge is enormous and there are many different foci of research. Our chapter concentrates on research with human adults. Readers interested in conceptual development or concepts in animals are referred to the reviews by Smith & Colunga (in press) and Lazareva & Wasserman (2008), respectively. For reasons of space, we must also ignore research on the impact of conceptual knowledge on inductive reasoning (for a review, see Hayes, Heit & Swendsen 2010) and formal models of categorization (see Pothos & Wills 2011).

2. Conceptual knowledge.

2.1. Functions of conceptual knowledge.

The most obvious function of conceptual knowledge is to allow for the categorization of things. In fact, this function has traditionally been in the focus of theoretical and empirical research on conceptual knowledge. It is commonly assumed that conceptual knowledge comprises distinct concepts, each of which provides information about a particular category of entities in the world (or more precisely, of entities that people deem as being in the external world). For example, the concept HAMMER may include information about what members of the category {hammers} look like, how they are used, and so on. (We indicate concepts by small caps and sets of entities in the world by curly brackets). A given thing is categorized by examining how well its properties match the information contained in a particular concept, possibly compared with alternative concepts. Theories differ in their assumptions as to the information contained in concepts and categorization decision rules (see Sec. 3 and 4).

A deeper understanding of what conceptual knowledge is good for is gained by considering its influence on representations of things in working memory (cf. Sec. 1). What happens when some part of a scene is recognized as a particular kind of entity, say, as a car? In what way does this use of conceptual knowledge shape the interpretation of the sensory input or, in other words, the mental representation of this part of the scene? Concept theories are not always explicit with respect to this issue but a widespread assumption seems to be that if something is conceptualized as a member of a particular category (e.g., {cars}), then its representation in working memory is essentially a replica of the content of the respective concept. This implies that whenever a person identifies things as a car, the working memory representations of those things are identical in content. Some more recent accounts ascribe greater flexibility to working memory representations. For example, simulation theory (Barsalou 2009) emphasizes that their contents are also influenced by the situational context. On this view, a car may be represented rather differently in working memory depending on whether it is being driven, filled with gas, washed, or bought. In any case, the accounts agree that conceptual knowledge affects representations in working memory in two complementary ways. On the one hand, some pieces of information are suppressed or deleted from the representation, specifically ones that are conceptually irrelevant. This may be considered the *abstraction* function of conceptual knowledge. Instead of representing the given thing in all its details, the representation mainly contains information that characterizes it as a particular kind of entity (in a particular situation). Abstraction is advantageous if not necessary to protect subsequent processing (e.g., thinking, problem solving, action planning) from being influenced by irrelevant information. On the other hand, the representation is supplemented with some pieces of information which stem from conceptual knowledge rather than being given by the stimulus itself. We refer to this as *prediction*. Prediction is a less obvious function of conceptual knowledge than abstraction. Let us therefore consider it in some more detail.

Conceptual knowledge is constantly used for predictions in daily life. When we grasp a hammer, we anticipate its approximate weight, even if we've never seen it before. When we cut an apple we expect it to be white inside. When we see a snowman in a backyard,

we assume it to have been built by children. Notice that the predictions considered here are not predictions in the ordinary sense. They may not only concern the future but also the present (e.g., the snowman is made out of snow) and the past (e.g., the snowman was built by children). Furthermore, they may be made unconsciously, and they derive from stored information about past situations rather than from explicitly learned rules. What is the basis for such predictions? Let us assume that in the current situation, there is something possessing the feature A. We further assume that according to the information represented in conceptual knowledge, previous situations with feature A also involved feature C, say, in 70% of the cases. Unless feature C's presence in the current situation is obvious anyhow, this knowledge can be used to estimate that the likelihood of feature C in the current situation is .70. Clearly, using feature A for estimating the likelihood of feature C is pointless if according to prior knowledge, C was present in 70% of *all* past situations (i.e., if according to prior knowledge, the base rate of C is .70). In this case, one could have estimated the likelihood of C to be .70 without considering feature A. However, taking into account feature A is advantageous if according to prior knowledge, feature C was more often, or alternatively, less often present in situations containing A than in other situations. More generally speaking, taking into account a given feature A improves the prediction of a yet unobserved feature C, if there is a statistical **association** between the features A and C. Of course, usually more than a single feature A is used to estimate the likelihood of an unobserved feature C, and other features or feature combinations, say B, may modify the association between the features A and C. For example, the likelihood of feature C <breaks when dropped> is high for an object with feature A <cup-shaped> if feature B <made of porcelain> is present but low if feature B' <made of plastic> is present instead.

It should be noted that in the literature on concepts and categorization, one frequently finds the term *correlation* instead of *association*. However, as features are usually considered qualitative properties (i.e., being either present or absent), it is most often the contingency between two features that is at issue. We therefore use the umbrella term *association* to cover both correlation and contingency.

In sum, feature prediction uses information about associations among features in past situations. This information is provided by conceptual knowledge. As we have seen, conceptual knowledge need not be organized into distinct concepts to allow for feature prediction. However, as mentioned, many accounts postulate such an organization. These accounts generally consider categorization a necessary first step for feature prediction. Specifically, a given thing is first assigned to a particular category and then the information contained in the respective concept becomes available (see Murphy & Ross 2010 and Hayes, Heit & Swendsen 2010: 286–287, for a discussion of this issue).

2.2. The content of conceptual knowledge.

Our characterization of conceptual knowledge as knowledge about feature associations converges with a view that has been widespread since Rosch's seminal articles (e.g., Rosch 1978; Rosch et al. 1976). On this view, conceptual knowledge has its basis in the correlational structure of the world. Each concept represents a particular bundle of strongly associated features (cf. Sec. 3 for other views). Of course, the features are not features of the world as such but arise from our sensory-motor and emotional systems (e.g., <red>

<sticky>, <ugly>) and higher cognitive processes integrating information from various sources (e.g., <dangerous>, <breakable>, <expensive>).

The relationship between features and concepts is an intricate matter. First, at least many high-level features arise so as to facilitate the discrimination of categories (see Schyns, Goldstone & Thibaut 1998). Thus, rather than being independent building blocks of concepts, features themselves may to some extent depend on required conceptual distinctions. Second, high-level features are probably often configurations of simpler features. From a structural point of view, such features are therefore difficult to distinguish from concepts. Moreover, features may even involve concepts (e.g., <has a pit>, <eats meat>). This entails a significant broadening of the notion of features, and in addition, it introduces a new aspect of conceptual structure, namely that of thematic relations. Let us briefly explain this issue.

In Section 2.2 we deliberately spoke of features *in situations*. People usually do not experience isolated things. Rather they experience things in the context of particular situations and as objects of their own actions. It is likely that associations between the features of a given thing (e.g., a cherry) and features of things frequently encountered in its context (e.g., tree), as well as features of actions frequently performed with the thing (e.g., picking, eating) are encoded. Thus, conceptual knowledge also contains information about so-called thematic relations (e.g., cherry – tree; cherry – eating; hammer – nail; for empirical evidence, see, e.g., Estes, Golonka & Jones 2011). Having a particular thematic relation to other entities can be considered a feature of an entity. Such features are sometimes called *extrinsic* features as opposed to *intrinsic* features, which are true of an entity in isolation (see Barr & Caplan 1987). Thus, for example, the concept CHERRY may include not only the intrinsic features <red> and <round> but also the extrinsic features <grows on trees>, <can be bought on the market>, <can be eaten>. Indeed, many studies have shown that commonalities with respect to extrinsic features increase perceived similarity and affect categorization (for a review, see Estes, Golonka & Jones 2011).

We have emphasized the correlational structure of the world as the basis of conceptual knowledge. However, subjective factors play an important role as well. It is reasonable to assume that the feature associations that get encoded are mainly those that are sufficiently salient and relevant to a person's life. Thus, cultural background, job, and interests may have a significant impact on people's conceptual structure (see, e.g., Medin et al. 2006, and Tanaka & Taylor 1991). It may even be the case that only feature associations construed as causal relations are encoded (see Sec. 4.3 for a discussion). Another important issue is that not all feature associations encoded in conceptual knowledge stem from direct experience; many of them may derive from communication with other people.

2.3. Conceptual hierarchies.

Things can often be categorized in various ways. For example, something may be conceptualized as a flute, a component of an orchestra, a gift for a child, a thing to take on a vacation, a recorder, or a musical instrument. Particularly the possibility of identifying things at various levels of specificity (e.g., musical instrument, flute, recorder) has received much attention in research on concepts and categorization. Which level is preferred and why? Before addressing this question, let us consider the conditions for differentiating a concept into more specific concepts.

Establishing concepts at a more specific level of abstraction is not done arbitrarily, but according to certain constraints. For example, most people lack concepts for different types of mountains or ideas. Furthermore, while `RED_WINE` and `WHITE_WINE` are well-established concepts, `RED_DRESS` and `WHITE_DRESS` are not. Why not? If we take into account that a main function of concepts is feature prediction, the answer is straightforward. A concept such as `RED_DRESS` would not allow for any predictions other than those inferable from the concept `DRESS` plus the information that the dress is red. In contrast, the concept `RED_WINE` allows additional predictions with respect to the taste of the wine, its optimal temperature, and the meals that it goes well with. Such predictions are possible because in the category {wines}, certain colors are associated with certain tastes, optimal temperatures, and appropriate meals. In other words, our conjecture is that a more specific concept is established only if there are feature associations *within* the category specified by the parent concept. The more specific concept then renders it possible to predict new features that cannot be predicted on the basis of the parent concept. It may be interesting to note that according to this view, it is unlikely that the concept `BACHELOR` only comprises the features listed in the concept `MAN` (<human>, <male>, <adult>) plus the feature <unmarried>. If conceptually a bachelor were no more than an unmarried man, then the concept would not exist. Rather, the concept captures the association between <unmarried> and certain other features occurring in the category {men}, as for instance, <has to take care of the laundry himself>, <is not responsible for a family>, <is always ready to go to a party>, and so on.

Based on these considerations, it may be supposed that people prefer using concepts at the lowest level, since that allows the most predictions. In their renowned study, Rosch et al. (1976), however, demonstrated that the level that people prefer in conceptual tasks (dubbed the *basic level*) is most often a certain middle level in a taxonomy. For example, people prefer categorizing things as members of {chairs}, {tables}, or {beds} rather than as members of {kitchen chairs} or {pieces of furniture}, and similarly, they prefer using the categories {flutes}, {drums}, {pianos} rather than the subordinate categories (e.g., {recorders}) or the superordinate category {musical instruments}. Many subsequent studies replicated this finding, and in addition provided evidence that basic level superiority is not simply due to the fact that the labels of basic level categories are relatively frequent and short words and are acquired relatively early in childhood (for reviews, see Mervis & Rosch 1981 and Murphy & Lassaline 1997). It should be noted, however, that the basic level is not always privileged. For experts in a domain (e.g., dog experts, bird watchers), the subordinate level is as useful as the basic level in their domain of expertise (e.g., `BEAGLE` or `COLLIE` vs. `DOG`; see Johnson & Mervis 1997; Tanaka & Taylor 1991), and in semantic dementia, the superordinate level appears to be better preserved than the basic level (see Rogers & Patterson 2007).

Why is the basic level usually privileged in conceptual tasks? Important hints come from studies in which participants were asked to list as many features as possible that are shared by the members of a given superordinate, basic, or subordinate category, respectively (e.g., Johnson & Mervis 1997; Rosch et al. 1976; Tanaka & Taylor 1991). Not surprisingly, participants listed more features for categories lower in the hierarchy (e.g., the number of listed features increases from {pieces of furniture} to {chairs} to {kitchen chairs}). Yet, the increase was not constant for each downward move but largest when moving from the superordinate level (e.g., {pieces of furniture}) to the basic level (e.g., {chairs}). The move from the basic to the subordinate level (e.g., {kitchen

chairs)) yielded relatively few additional features. This suggests that basic-level categories are much more homogenous than superordinate categories, which is clearly advantageous with regard to category-based feature prediction. In addition, Rosch et al. (1976) found that the members of a basic category (e.g., {chairs}) share relatively few features with the members of other categories at the same level, (e.g., {tables}, {beds}, {cupboards}). In other words, alternative categories at the basic level are particularly clearly differentiated from each other, compared with alternative categories at other levels (e.g., at the subordinate level: {kitchen chairs} vs. {office chairs} vs. {easy chairs}).

On the basis of these findings, Mervis & Rosch (1981) characterized the basic level as the level at which the set of entities of a domain is partitioned in such a way that the categories maximize within-category similarity relative to between-category similarity. Other researchers have also been concerned with the *structural* properties of the partitions at different levels in natural and artificial taxonomies and with possible measures of the utility of partitions (for an overview and a recent proposal, see Gosselin & Schyns 2001).

Another possible explanation of the privileged status of basic-level categories emerges from the consideration that outside of the laboratory, the different kinds of features typically play different roles in conceptual processing. For example, what is typically “given” when perceiving things are salient visual features (e.g., shape, part structure, color, movement), whereas the features that we want to predict are the features that arise over time (e.g., the melting of a snowman), the appropriate motor programs for interacting with the thing, and the features that emerge from this interaction (e.g., the weight of a hammer; the behavior of a rabbit when one approaches it), as well as more abstract features. Thus, concepts should capture the associations between visual features and these latter kinds of features. Superordinate concepts may be largely useless in this regard. Rosch et al. (1976) (see also Jolicoeur, Gluck & Kosslyn 1984; Tanaka & Taylor 1991) found that different superordinate categories (e.g., {vehicles}, {buildings}) can hardly be distinguished on the basis of visual properties; they mainly differ in function (e.g., <used for transport of persons>). In contrast, basic-level categories (e.g., {cars}, {trucks}, {airplanes}) were found to typically differ in salient visual features (shape, part structure) as well as with respect to motor movements for interacting with the things. Thus, identifying a thing as an instance of a particular basic-level concept is probably relatively straightforward and allows predictions as to appropriate interactions. The subordinate level may again be less useful. At this level, the information about the visual features is refined but with respect to action affordances and appropriate motor programs there are no significant differences between the different subordinate categories of the same basic category (e.g., {Rolls Royces} vs. {Mini Coopers}). In other words, categorization at the subordinate level costs more perceptual effort (see, e.g., Collin & McMullen 2005) without providing more information as to how the given thing can be interacted with (clearly, there may be a profit with respect to the prediction of some other features, e.g., <expensive>). Taken together, it is plausible that when perceiving things, the natural “entry point” into conceptual knowledge is at the basic level (Jolicoeur, Gluck & Kosslyn 1984). Let us add, however, that this probably does not apply to all conceptual tasks. For example, in action planning, when pondering about possible means of achieving a particular goal, functional features may come to mind first. Someone who is hungry may think of buying something that can be eaten, and someone planning to go to a dangerous place may think of taking along something for

self-defense. Thus, in action planning, the entry point into a taxonomy may often be at the superordinate level.

We end this section with a cautionary remark on the notion of hierarchical relations. The organizational principle of a truly hierarchical classification scheme is that of set inclusion. However, it is questionable whether this principle generally applies to concept-based taxonomies. First, concepts, as we have characterized them, do not provide defining features of the members of a category but features are more or less biased to the assignment of a given entity to a particular category (cf. Sec. 3 and 4). Hence, intransitive categorical decisions may arise. For example, a car seat may be judged to belong to the category {chairs} but not to the category {pieces of furniture} (Hampton 1982; see also Sloman 1998). Second, many of the concepts that are commonly considered superordinate concepts (e.g., CLOTHING, FOOD, JEWELRY) may actually refer to groups or collections of heterogeneous entities, united by spatio-temporal contiguity and function (see Wisniewski, Imai & Casey 1996). If so, then a single item, for example a shirt, can not more be considered an instance of CLOTHING, than a single singer an instance of CHOIR or a single ship an instance of FLEET.

3. Theoretical approaches.

Although in the previous section we tried to avoid committing ourselves to a particular concept theory, our presentation was certainly not theory-neutral. This section gives an overview of the theoretical approaches to conceptual knowledge and categorization.

According to the **definitional approach** (or classical approach), a concept defines a category by specifying the features that are singly necessary and jointly sufficient for membership in the category. Few if any cognitive psychologists consider this view adequate for everyday concepts and categories. This is not to deny that in certain kinds of artificial or technical category learning tasks, people may expect the categories to be well-defined and aim at finding a simple rule for discriminating between them (*rule-based categorization*; see Close et al. 2010). However, there are a number of strong theoretical and empirical arguments against the definitional view of everyday concepts (see, e.g., Murphy 2002, chap. 2). Many of them were pointed out by Rosch and her colleagues in a series of seminal articles in the 1970s (for an overview, see Mervis & Rosch 1981). These researchers also proposed an alternative to the definitional view which quickly found many adherents, namely the prototype view (e.g., Rosch & Mervis 1975).

According to the **prototype view**, the different members of a category, rather than all sharing a certain set of features, each match (or resemble) other members in different respects. In other words, they bear a “family resemblance” (Rosch & Mervis 1975). Category membership is a matter of degree; it is a function of an item’s similarity to the *prototype* of the category, which is what is represented in the corresponding concept. There are two rather different conceptions of a prototype. According to the first one, it is an assemblage of all possible features, each weighted by its frequency of occurrence in the category (e.g., Rosch & Mervis 1975) or by another measure of its importance for the category (e.g., Hampton 1993). For example, in the prototype for TOMATO, <red> has a greater weight than <green>. According to the second conception, a prototype is a sort of central-tendency instantiation of the category, possessing the features that correspond to the mean or modal value of the category members on each attribute dimension

(e.g., Minda & Smith 2011). Notice that neither conception envisages that a prototype captures within-category relations between attribute dimensions or the relative frequency of co-occurrence of certain features in the category (e.g., <red> & <ripe>). Rather, the various attribute dimensions are considered independently of each other. Hence, prototype theories belong to the class of *independent cue theories* (Medin & Schaffer 1978).

The previously mentioned theoretical approaches regard concepts as knowledge structures that – albeit possibly being used as building blocks in other types of knowledge – are in principle independent of other types of knowledge. By contrast, the **theory-based approach** (sometimes referred to as explanation-based or knowledge-based approach) assumes that concepts are embedded in naïve domain-specific theories (e.g., Murphy & Medin 1985). Concepts are “mini-theories” (Rips 1995), specifying categories in terms of causal relationships among features. Category membership is determined by estimating how well the features of a given thing can be explained by the causal mechanisms specified in the concept (e.g., Rehder 2010). One version of the theory-based approach is **psychological essentialism** (e.g., Gelman 2004; Medin & Ortony 1989), according to which people believe that the members of a category share an unchanging property, an essence, that causes category members to have the features they do. The essence of a category may be unknown, in which case the concept contains an “essence placeholder”. Notice that psychological essentialism, like the definitional approach, assumes that categories have clear-cut boundaries – every entity either is or is not in a particular category. This does not imply clear-cut categorization judgments. Often a person may be uncertain about the essence of a given thing and needs to rely on features considered diagnostic of essences.

Almost all theories posit that a concept is a sort of summary representation of a category, characterizing the set of category members as a whole. The only exception is the **exemplar approach**, which assumes that a concept represents the individual exemplars of the category that have been encountered in the past (e.g., Medin & Schaffer 1978; Nosofsky 1986; Storms 2004). The “glue” holding together the different exemplars of a category is their common label. A thing with an unknown label is categorized by comparing it with the individual exemplars of the relevant alternative categories and choosing the category for which the observed similarities are largest overall. Exemplar theories imply that people possess implicit knowledge about the co-occurrence of features within categories and that categorization is sensitive to the particular combination of features being true of the given thing.

According to **connectionist models**, conceptual knowledge is encoded in a large network of representational units with weighted connections between them. In distributed models (e.g., McRae 2004; Moss, Tyler & Taylor 2007) the units represent conceptual microfeatures and the weights of the connections reflect the strengths of their associations. When a group of microfeatures becomes activated (e.g., by sensory input), activation is propagated through the network via the connections until eventually a stable pattern of activated units is reached. This pattern is a working memory representation that is shaped by conceptual knowledge. However, the process does not necessarily imply categorization in the usual sense, as distributed connectionist models do not generally assume the conceptual network to be organized into distinct concepts.

A related theory is the **simulation view** of conceptual processing (e.g., Barsalou 2009), which assumes that concepts are bindings of memory traces distributed over

modality-specific mental subsystems. Importantly, concepts include information about the situations in which the category's members were encountered. Upon perceiving an entity, its features and the context entail a re-enactment of various memory traces that were formed when similar things were previously encountered in similar situations. The result is a highly situation-specific construal of the given thing as a member of the category.

The various theoretical views are not mutually exclusive. For instance, a prototype model may make the additional assumption that concepts contain information about causal and other relations between features (see Hampton 2006). Furthermore, summary representations and sets-of-exemplars representations are frequently taken as end points of a continuum. Example models include Anderson's (1991) rational model, SUSTAIN (Love, Medin & Gureckis 2004), and the varying abstraction model (Vanpaemel & Storms 2008). Other models assume that people draw on different kinds of knowledge when categorizing items, for instance, on prototypes plus remembered exemplars (e.g., Smith & Minda 2000). Similarly, some researchers emphasize that multiple, neurobiologically distinct memory systems contribute to category learning and categorizing (e.g., Ashby & Maddox 2011).

4. Issues of debate.

Many controversies in research on concepts and categorization originate from criticisms of the prototype view. In this section, we address three important issues of debate. In each of them certain implications of the prototype view are compared against those of one or two other theoretical views mentioned in the previous section.

4.1. Is category membership a matter of degree?

Prototype theories assume that category membership is graded. The more similar a given thing is to the prototype the more clearly it is a member of this category. In addition, according to prototype theories, the typicality of an item reflects its degree of category membership. Both these assumptions are questioned by other researchers, in particular by proponents of the definitional and essentialist view. They posit that category membership is all-or-none – a thing is either a full member of a category or it is not a member of the category – and typicality has nothing to do with category membership. Let us first consider the variable of typicality and then turn to the more general question of whether category membership is all-or-none or a matter of degree. Before reviewing the empirical findings, it is important to re-emphasize that in empirical research on conceptual knowledge, the categories and the items to be categorized are often specified linguistically, and this is especially true in this research area. Most of the findings we report in this section are therefore actually findings about knowledge and use of lexical meanings. However, we present them in accordance with the way they are normally interpreted.

It is well-established that members of a category vary in the degree to which they are considered representative or good examples of the category. One particular cat may appear “cattier” than another one. Similarly, a trout or a herring is considered more representative of the category {fish} than a shark or a flounder, for instance. The most common measure of representativeness is typicality. Typicality is operationally defined, namely by responses to questions of the form *How typical is item x of category y?* or *How good an*

example is item x of category y? Notice that typicality is a matter of the *relation* between an item and a category. This becomes evident when we consider different levels of a taxonomy. For example, a robin is rated more typical than a chicken if the target category is {birds}, but the opposite is true if the target category is {animals} (see Roth & Mervis 1983).

It should be mentioned that ratings of typicality are not always based on considerations concerning representativeness. Specifically, in domains in which a person has expert knowledge, and with goal-derived categories (e.g., {foods to eat on a diet}), the ratings are mainly determined by how close an item is to the ideal of the category (see, e.g., Barsalou 1985; Lynch, Coley & Medin 2000; for a unified account, see Davis & Love 2010). However, we ignore this “atypical” variant of typicality in the following.

Empirical research has revealed that typicality plays a role in a wide variety of conceptual tasks (for reviews, see Mervis & Rosch 1981 and Smith & Medin 1981, chap. 3), as well as in lexical processing and the pragmatics of certain expressions (see, e.g., Onishi, Murphy & Bock 2008; Rosch, 1978). Most importantly in our context, typicality has been found to be highly correlated with category-membership judgments as well as with measures of feature overlap and other measures of similarity to the prototype (e.g., Hampton 1998; Rosch & Mervis 1975). Proponents of the prototype view consider these findings as evidence that typicality is based on the same underlying variable that category membership is based on, namely similarity to the prototype (see Hampton 2007 for an explication of this assumption). This conclusion is challenged by other researchers (e.g., Armstrong, Gleitman & Gleitman 1983; Kamp & Partee 1995; Osherson & Smith 1997), who argue that for theoretical reasons and in view of certain empirical findings, typicality and category membership need to be distinguished. For example, Armstrong, Gleitman & Gleitman (1983) point out that graded typicality judgments are obtained even for well-defined categories such as {even numbers}. A summary of the main arguments in this debate is given in Hampton’s (2007) rejoinder.

Considering the controversial status of typicality ratings, it is reasonable to ask participants directly for judgments of category membership to find out whether category membership is absolute or a matter of degree. The simplest and most frequently used method of obtaining category-membership judgments are Yes-No categorization tasks. Participants are presented with a category name and various items (pictures of objects or verbal labels) and are asked to decide for each item whether or not it is a member of the category. Typically, a gradient of judged category membership is found – some items are categorized as members of the target category by more participants than are others. For example, in a study by McCloskey & Glucksberg (1978), the item *airplane* was categorized as a member of the category *vehicles* by nearly all participants whereas *roller skate* and *parachute* turned out to be “borderline” items, judged as members of the category *vehicles* by barely more than 50% of the participants. Of course, this finding may simply reflect individual differences in the placement of the category boundaries. However, McCloskey & Glucksberg (1978) also found that participants, when presented with the task a second time, sometimes changed their categorization decision, in particular for borderline items. This variability may be attributed to an instability of the criteria for judging category membership or to fluctuations in the content of the representations established in working memory. In any case, the finding suggests that judged category membership is more “fragile” for some items than for others. This however does not yet prove that people believe that category membership is a matter of degree.

To clarify this issue, other experimental paradigms were developed. For example, participants were asked to judge the category membership of items on a scale from “definitely not a member” to “definitely a member”, offering the opportunity for expressing degrees of category membership (e.g., *How clearly is an escalator a member of the category ‘vehicles’?*) (see, e.g., Barr & Caplan 1987; Diesendruck & Gelman 1999). Furthermore, various meta-cognitive tasks were used. For example, Kalish (1995) presented pairs of statements such as *John says this animal is an elephant* and *Jane says this animal is not an elephant* and asked participants to decide whether this disagreement would in principle be resolvable as only one statement can be true, or whether it would in principle be irresolvable as one can always argue for both sides. Together, the results from these studies (see Estes 2004 and the literature cited therein) suggest that people consider category membership a matter of degree for artifacts (e.g., vehicles, tools), while they are somewhat more inclined to assume absolute membership for many categories of natural kinds (e.g., birds, fruit). It should be added that differences between concepts of artifacts and natural kinds have been revealed in other areas of research as well, but there is as yet no widely accepted answer as to what precisely distinguishes the concepts in these domains (see Margolis & Laurence 2007).

4.2. Summary representations or representations of sets of individual exemplars?

Much research has been devoted to the question of whether concepts provide information about entire categories (“summary representation”) or represent individual exemplars of categories. Although the former view is taken by many different theories, the debate is centered between those versions of prototype theory that consider prototypes as central-tendency representations (see Sec. 3) and exemplar theories. The debate led to a flood of categorization studies, mostly using artificial categories (e.g., sets of dot patterns; sets of geometric forms varying in shape, size, and color) that participants first learn in the experiment. Using artificial stimuli has the advantage that the categories can be tailored to the question at hand. To illustrate, let us consider a simple categorization task in which participants assign stimuli to one of two categories, {a} and {b}. According to prototype theory, all that matters is the similarity of the given stimulus *S* to the prototype of each of the categories (see Minda & Smith 2011). The more similar *S* is to one of the prototypes, the more likely it will be categorized as a member of this category. Specifically, if *S* exactly matches one of the prototypes, say that of category {a}, then the likelihood that it is categorized as a member of category {a} rather than {b} is maximal, even if it is quite similar to some exemplars in the alternative category. According to exemplar accounts, however, what matters is the similarity of *S* to the individual exemplars of the two categories (see Nosofsky 1986). Even if *S* is identical to the prototype of one of the categories, it may be categorized as a member of the contrast category, provided it is extremely similar to one or more of the exemplars of this category. By creating artificial categories, variables such as these can be manipulated, while keeping other ones constant. Importantly, whereas in many natural categories the prototype is an abstract entity, which doesn’t actually exist, artificial categories can be designed such that the prototype exists in the set of stimuli.

To test the validity of the theories, many studies have investigated the categorization of the prototypes of categories. In a typical experiment, the stimulus material comprises two stimuli constituting the prototypes of two categories (e.g., two different patterns of five dots each) and a number of different “distortions” of the prototypes (e.g., patterns of five dots that slightly differ from the respective prototype with respect to the spatial relations among the dots). In a training phase, participants are presented with a selection of the distortions and are told which category each pattern belongs to. The prototypes themselves are not presented in this phase. In a later transfer phase, participants categorize old distortions (i.e., patterns that were presented during training), new distortions (i.e., patterns that were not presented before), as well as the prototypes of the categories. According to prototype theory, the prototype of a category should be particularly easy to categorize (*prototype-enhancement effect*). Early studies using this prototype-distortion paradigm (e.g., Posner & Keele 1968) confirmed this prediction. However, in those studies, the similarity between the prototype of a given category and the old distortions belonging to the same category was on average higher than the similarity between the new distortions and the old distortions in the same category. For such a situation even exemplar models predict a prototype-enhancement effect (see, e.g., Shin & Nosofsky 1992). More recent studies that tease apart the relevant variables (i.e., similarity to the prototype vs. similarity to other exemplars of the categories) support the predictions of the exemplar view (cf. Nosofsky 2000, but see Minda & Smith 2002).

A possible drawback of artificial categories is that researchers inadvertently create conditions that favor one or the other account. For example, if the experiment involves only a few categories, with a small number of exemplars per category and little “within-category structure”, then it isn’t surprising that participants tend to encode and remember the individual exemplars. In contrast, if multiple, large and highly structured categories are to be learnt then creating summary representations may be advantageous. Indeed, the results mentioned above that favored exemplar models mostly stemmed from studies that employed only two small categories with little within-category structure. When conditions were less favorable for memorizing individual exemplars, results were more in line with prototype theories than with exemplar models (for an overview, see Minda & Smith 2011). To account for these findings, in recent years various hybrid categorization models have been proposed (see Sec. 3). In addition, increasingly more attention has been devoted to whether the findings generalize to natural language categories (see, e.g., Storms 2004) and to a wider range of category uses (see, e.g., Markman & Ross 2003).

4.3. Relations between features within a category.

Rosch (1978) emphasized that conceptual knowledge captures the correlational structure of the world. Surprisingly, however, according to Rosch’s and other prototype theories, feature associations *within* categories are not encoded in conceptual knowledge. To illustrate, let us consider the category {spoons} (see Medin & Shoben 1988): Spoons differ from one another with respect to the material they are made of and their size, among other things. For the category {spoons}, these two attribute dimensions are associated: Wooden spoons tend to be relatively large, whereas metal spoons are more often small or medium-sized. Notice that this is a *within-category* relation, which possibly only holds for {spoons}. Across the board, material and size may be unrelated, and in certain other

categories there may even be an association in the opposite direction (e.g., in the category {ships}). In any case, according to prototype theories, associations such as these are not captured in a concept. However, other theories do assume that concepts contain information about within-category featural associations. According to exemplar theories, the information about the statistical co-occurrence of features within a category is implicitly coded in the knowledge of the category's exemplars. Connectionist and theory-based accounts both posit explicit representations of feature associations but their assumptions differ in an important respect: Connectionist accounts imply that statistical co-occurrences are encoded; the more often two features are encountered together, the greater the weight of the connection between the respective units (see, e.g., McRae 2004). By contrast, theory-based accounts assume that mainly those feature relations are encoded for which the person has a causal explanation (e.g., Ahn et al. 2002; Murphy & Medin 1985). We will refer to these two kinds of associations as *statistically-based* and *theory-based* associations, respectively.

It is now well-established that, contrary to what prototype theories imply, conceptual knowledge does encode within-category feature associations. Much evidence comes from studies conducted in the theory-based framework. For example, it was found that the status of a feature in the structure of causal relations (supposed to be represented in a concept) affects how much importance is attached to this feature in categorization decisions (*causal status effect*). Furthermore, objects are classified by evaluating whether their features are likely to have been generated by the structure of causal relations that make up the concept (for an overview, see Rehder 2010).

However, these studies have been exclusively concerned with theory-based associations. It remains open whether purely statistically based feature associations are encoded in conceptual knowledge as well, as connectionist and exemplar accounts imply. A study by Ahn et al. (2002) suggests that the associations people are aware of are mostly ones they conceive of as causal relations. However, this conclusion was challenged by McNorgan et al. (2007). Moreover, it must be borne in mind that not all information encoded in conceptual knowledge is necessarily conscious. Indeed, various studies demonstrate that people possess and use knowledge of feature associations that they have probably never consciously thought of and that they may not be able to provide an explanation for (e.g., McNorgan et al. 2007; McRae 2004; Murphy & Ross 2010).

5. Conceptual combination.

Conceptual combination is the process by which a complex representation is constructed from two or more concepts. Using almost exclusively linguistic stimuli, research on conceptual combination is effectively concerned with the interpretation of complex linguistic expressions, mainly nominal expressions such as *brown apple*, *sports which are games*, or *mountain bird*. The result of conceptual combination is frequently called a *complex concept*. However, it is important to keep in mind that it is actually a representation in working memory, not a novel long-term memory structure (see Sec. 1). We refer to the result of conceptual combination as a *composite working memory representation* or simply *composite representation*.

After an initial debate about the viability of an extensional analysis of conceptual combination in the early 1980s (e.g., Osherson & Smith 1981), research has focused on

intensions, that is, on the properties represented in concepts and in composite representations constructed from them. Consequently, in empirical studies, participants are typically asked to generate or verify properties of the members of a named category or to describe their interpretation of a given stimulus expression in detail.

5.1. Empirical findings and theoretical approaches.

An early model of conceptual combination is the Selective Modification Model (Smith et al. 1988), which is concerned with **adjective-noun combination** (e.g., *brown apple*). Noun concepts (e.g., APPLE) are assumed to have a prototype structure comprising a list of relevant attribute dimensions (e.g., color, shape, taste), with a set of weighted values for each dimension (e.g., <red>, <green>, <brown> for the color dimension in the concept APPLE). To create a composite representation for an adjective-noun phrase such as *brown apple*, the relevant attribute dimension (color) is selected and the weight of the value <brown> is enhanced whereas the weights of the other values are set to zero. With a few additional assumptions, this model accounts for typicality-judgment phenomena observed with adjective-noun phrases. However, the model is limited to adjectives that unambiguously refer to one of the attribute dimensions listed in the noun concept. It cannot deal with multiattribute adjectives (e.g., *shriveled apple* – shape and texture), nor with subsecutive adjectives (e.g., *good apple* vs. *good coffee*), nor with privative adjectives (e.g., *fake apple*). Several researchers (e.g. Medin & Shoben 1988; Murphy 1988, 1990) take this failure to reflect a more fundamental problem of prototype theories, which, according to their view, severely underestimate the richness of information used in conceptual combination (for another objection, see Connolly et al. 2007). For example, Murphy (1988, 1990) emphasizes that world knowledge plays a pivotal role in conceptual combination, and Medin & Shoben (1988) call attention to the context dependence of adjectives. We shall come back to these arguments in Section 5.2. Since this debate, only few articles addressing adjective-noun combination have been published (e.g., Franks 1995).

Much more research has been devoted to **noun phrases containing a restrictive relative clause** (e.g., *sports which are games*), which are thought to require the conjunction of noun concepts for their interpretation. Research in this field is strongly influenced by Hampton's (1987) Composite Prototype Model, which posits prototype concepts representing property lists, where the properties are weighted according to their importance for the respective concept. A composite representation is formed by merging the properties of the two constituent concepts, assigning new weights to the properties according to certain rules, and performing a consistency checking procedure to ensure that the composite representation does not inherit incompatible properties. This model accounts for a wide range of empirical findings, including concept dominance, non-commutativity, and over-extension (see Hampton 1997). There is a phenomenon, however, that must be attributed to processes not captured by the model – the occurrence of *emergent properties*. We discuss this phenomenon in Section 5.2.

The current focus of conceptual-combination research is on (novel and familiar) **noun-noun compounds** (e.g., *mountain book*). There are two different theoretical approaches, the *schema-based* and the *relation-linking* approach. Schema-based theories (e.g., Costello & Keane 2001; Murphy 1988; Wisniewski 1997) posit rich concepts similar to those proposed by the theory-based view (see Sec. 3). Concepts are schema representations with complex internal structure, containing information about possible properties

and their interconnections as well as typical functional roles in scenarios (e.g., SPOON: instrument for eating, instrument for stirring), and other thematic relations to other concepts (see Sec. 2.2). According to these theories, a noun-noun compound is interpreted by integrating information from the modifier concept with the head concept. Wisniewski's (1997) Dual-Process Theory may serve as an example. Wisniewski distinguishes between different kinds of interpretations, the two most common ones being *relational interpretations* and *property interpretations*. These interpretations result from different processes. Relational interpretations occur if the modifier and head concept are found to fit different functional roles in a particular scenario. For example, *paint spoon* could be interpreted as a spoon used to stir paint because SPOON can be bound to the instrument role and PAINT to the object role in a stirring scenario. In contrast, property interpretations involve a mapping of one or more properties from the modifier concept onto the head concept. For example, *box clock* may be interpreted as a square clock and *zebra clam* as a striped clam. According to the Dual-Process Theory, property interpretations are the outcome of a process involving a comparison and alignment of the modifier and the head concept. However, recent findings suggest that the salience and diagnosticity of the modifier's properties play a crucial role (see, e.g., Costello & Keane 2001; Estes & Glucksberg 2000).

The relation-linking approach is inspired by traditional linguistic theories of compounding (see article 80 (Olsen) *Semantics of compounds*, Sec. 2) and was introduced into the research on conceptual combination by Gagné & Shoben (1997), see also Gagné & Spalding (2009). According to their theory CARIN (Competition Among Relations in Nominals), a compound noun is interpreted by linking the two constituent concepts via a thematic relation selected from a limited set including 'located', 'made_of', 'about', 'during', and some others. For example, *mountain bird* may be interpreted by selecting the relation 'located' (a bird in the mountains) and *mountain magazine* by selecting the relation 'about' (a magazine about mountains). To find a suitable interpretation, people exploit knowledge about statistical regularities in language use. More specifically, upon encountering a noun in modifier position, knowledge about its past use as a modifier becomes activated and the respective relations then compete for selection. CARIN predicts that, all else being equal, compounds instantiating a relation that has been used frequently with the given modifier are easier to interpret than compounds instantiating a relation less often used with this modifier. Thus, for example, as corpus analyses show that the modifier *mountain* is most often associated with the 'located' relation, *mountain bird* and *mountain tent* should be easier to interpret than *mountain magazine*. Empirical findings correspond to this prediction (e.g., Gagné & Shoben 1997; for left-headed compounds, see Storms & Wisniewski 2005).

It has often been questioned whether the great variety of relationships between nouns in compounds can indeed be reduced to a limited number of categories. However, let us accept this assumption and instead draw attention to CARIN's proposal concerning statistical knowledge. Considering the ample experimental evidence for the exploitation of statistical regularities in language comprehension (see Jurafsky 2003), the claim that people use statistical knowledge in processing compound nouns is no doubt plausible. However, does this knowledge actually concern individual lexical items, as Gagné & Shoben (1997) propose? Maguire, Wisniewski & Storms (2010) have reported that semantically similar words exhibit similar combination patterns in compounds.

With the benefit of hindsight, this finding is not surprising. Semantically similar words share many meaning components, and the way words are used in compounds certainly depends to a great deal on certain critical components. For example, the fact that *mountain*, when used as a modifier, is frequently associated with the relation ‘located’ but not with, say, the relation ‘during’ is most likely due to the fact that this noun denotes objects conceived as spatially extensive and permanent. Words that likewise denote objects conceived as spatially extensive and permanent (e.g., *valley, sea, city, garden*) can be expected to show the same preference for ‘located’ over ‘during’. More generally speaking, one may assume that certain critical meaning components (or combinations of them) are each associated with a characteristic relation frequency distribution. Thus, the statistical knowledge people use in processing compound nouns may actually concern relation frequencies associated with certain meaning components rather than lexical entries. Unlike Gagné & Shoben’s proposal, this revised proposal accounts for the finding that even rare words, which have probably never been encountered in compounds, display clear preferences for certain relations, when used as modifiers (Maguire & Cater 2005).

Certainly, the use of knowledge of statistical regularities can only be a *part* of the interpretation process. It only provides likely candidate relations. The respective interpretations must be worked out and their plausibility must be evaluated in order to eventually settle on one interpretation. CARIN says nothing about these processes. As several researchers (e.g., Storms & Wisniewski 2005) have pointed out, a fully-fledged theory should probably integrate assumptions from CARIN with those from schema-based theories. The processing of a compound noun may involve a statistic-based activation process (similar to the one assumed by CARIN or the revised proposal) as well as construction and evaluation processes (as proposed by schema-based theories).

As we have seen, research on conceptual combination has been dominated by the prototype view and the theory- or schema-based view. Connectionist theories and simulation theory are only recently becoming involved in the discussion (see, e.g., Levy & Gayler 2004; Wu & Barsalou 2009). The exemplar view, however, faces a particular problem. Its central assumption that a concept is a set of stored exemplars renders it difficult to account for the productivity of conceptual combination. Let us take the concepts STRIPED and APPLE as an example (see Osherson & Smith 1981). How can people form a composite representation out of these concepts if they have never encountered an exemplar of the category {striped apples}? Some authors (e.g., Storms et al. 1993; see also Storms 2004) propose to take representations of subcategories as the stored exemplars of a concept (e.g., PIGEON, CHICKEN, RAVEN are three exemplars of BIRD). However, it is as yet unclear whether this conception of exemplars solves the above-mentioned fundamental problem of the exemplar approach.

5.2. Conceptual compositionality.

Psychological research on conceptual combination, in which theories primarily seek to make correct predictions concerning people’s interpretations of complex expressions, has devoted relatively little attention to the issue of compositionality (for compositionality in semantics, see article 6 (Pagin & Westerståhl) *Compositionality*). Definitions of conceptual compositionality are most often left implicit, but many researchers would probably

agree that the composite representations created by means of conceptual combination are compositional to the extent that they are the result of processes that operate on the constituent concepts, guided by syntactic information. According to this definition, any use of information beyond that contained in the constituent concepts counts as non-compositional – even if this information is used in a rule-governed, predictable way. Hence, a non-compositional account of a given phenomenon is not necessarily inferior to a compositional account as far as predictive power is concerned. It should also be borne in mind that composite representations are representations in working memory, constructed in particular situations. Every reasonable theory must therefore acknowledge that at some point in time during the construction of a composite representation, information external to the constituent concepts is used. Of course, whether or not a theory provides a compositional account of a finding is still an important question. Two sets of findings have received particular attention in the discussion, context-dependent property instantiation and emergent properties.

Several researchers have pointed out that properties (even ones named by intersective adjectives) are instantiated differently in composite representations depending on the entities they are applied to (e.g., Medin & Shoben 1988; Murphy 1988; Wisniewski 1997). For example, when *zebra clam* is interpreted as a black-and-white striped clam, the stripes on the clam are probably represented as smaller and thinner than those of a zebra (Wisniewski 1997; see also Wilkenfeld & Ward 2001). Context dependence is often taken as evidence that the properties in a concept are not independent of one another, so that a new property must be accommodated to the particular selection of properties in the head concept (but see Rips 1995 for a more differentiated view). If this is true, then theory-based or schema-based theories are, in principle, able to provide a compositional account, while prototype theories are not, as a prototype concept contains no information about the interrelations among properties (see Sec. 3 and 4.3).

As we mentioned above, conceptual combination sometimes yields emergent properties (or *phrase features*): People consider certain properties true or typical of the composite but not of either of its constituents. For example, Hampton's (1987) participants often listed the property <live in cages> for *pets which are birds* but did not judge this property true of pets nor of birds. Emergent properties have also been observed with other kinds of combinations (e.g., *large spoon* → <wooden>; *beach bicycle* → <equipped with wide tires>) (see, e.g., Hampton 1997; Medin & Shoben 1988; Murphy 1988; Wilkenfeld & Ward 2001). It is widely agreed that there are at least two possible sources of emergent properties. First, they may arise from *extensional feedback*. That is, if the complex expression denotes a category of familiar things (e.g., *large spoons*; *pets that are also birds*), then people may retrieve familiar instances of this category from long-term memory and “look up” their properties. Notice that according to prototype and theory- or schema-based theories, extensional feedback can only take place *after* the combination of the concepts – the combination process cannot peek at instances of the composite (*no-peeking principle*, Rips 1995). Hence, these theories cannot provide a compositional account of emergent properties arising from extensional feedback. In line with the notion of extensional feedback, empirical studies (e.g., Swinney et al. 2007) have shown that emergent properties become available at a later point in time during the processing of a complex nominal expression than do “inherited” properties. For example, when presented with *peeled*

apple, people need more time to access the emergent property <white> than the inherited property <round>.

Second, emergent properties may result from reasoning processes and domain theories. An example is when *helicopter blanket*, interpreted as a cover for a helicopter, is ascribed the emergent property <waterproof> (Wilkenfeld & Ward 2001). Obviously, prototype theories must attribute these properties to reasoning processes taking place after concept combination. In contrast, theory-based approaches may be well-suited to provide a compositional account, by explaining how these properties emerge from reasoning processes that use the mini-theories contained in the concepts. However, as yet there are no systematic studies on this issue.

6. Relationship between conceptual knowledge and word meanings.

As we mentioned in the previous sections, many authors assume that conceptual knowledge is organized into distinct packages of information, that is, into concepts, and take concepts to be equivalent to lexical meanings. In the present section, we consider the implications of this equivalence view, and describe possible alternatives. The issue has much in common with the dictionary/encyclopedia debate in linguistics (for an overview, see Peeters 2000; see also article 30 (Jackendoff) *Conceptual Semantics*, and article 31 (Lang & Maienborn) *Two-level Semantics*). However, instead of repeating arguments from this debate, we look at the issue from a different perspective, examining whether the equivalence view is compatible with theoretical considerations and empirical findings from research on conceptual knowledge. Accordingly, the focus is on the relationship between concepts and meanings of count nouns. We disregard morphological issues and up until the end we also ignore homonymy, polysemy, and synonymy. Thus, we simply speak of *word meanings* and assume that each of them is associated with exactly one *word form*, and vice versa.

6.1. Equivalence view.

The view that concepts are equivalent to the meanings of words comes in two variants (see Fig. 106.1). On the first variant, which seems to be the prevalent one in research on concepts in cognitive psychology, concepts *are* word meanings. On the second variant, concepts and word meanings are stored in different mental subsystems but for each concept there is exactly one word meaning that shares the same informational content and vice versa. This variant is more in line with the traditional notion of a mental lexicon as a mental subsystem clearly separated from non-linguistic long-term memory.




Fig. 106.1: Two variants of the equivalence view. Left: Concepts are word meanings. Right: Concepts and word meanings are distinct but informationally equivalent knowledge structures. C_i = concept, M_i = word meaning, F_i = word form.

One problem with the equivalence view is that it implies that all concepts have a label. Each concept is connected with one particular form, either directly (Variant 1) or indirectly via the corresponding meaning (Variant 2). This entails a very narrow perspective on conceptual knowledge, which disregards concepts in pre-verbal infants and animals (for reviews, see Rakison & Yermolayeva 2010 and Lazareva & Wasserman 2008, respectively) and also a good deal of conceptual knowledge of human adults. Concepts can no doubt be acquired incidentally and without learning labels for them. This has been shown experimentally by studies of unsupervised category learning (e.g., Billman & Knutson 1996) and is also evident in our everyday experience. For example, people often discriminate between several subcategories of things (e.g., different subcategories of trees) without having names for them (see Malt & Sloman 2007: 102), and they possess concepts for which there is no label in their language (e.g., `EXTENDING_ONE'S_ELBOWS_TO_THE_SIDE_TO_FEND_OFF_A_SHOVING_NEIGHBOR`; see Murphy 2002: 389).

A second important implication of the equivalence view concerns the issue of the universality of the structure of conceptual knowledge. Different languages carve up the world in different ways, and, according to the equivalence view, conceptual distinctions match lexically coded distinctions. Speakers of different languages must therefore be assumed to differ in the structure of their conceptual knowledge. In recent years, there has been a revival of interest in the Sapir-Whorf or linguistic relativity hypothesis, which states that language influences thought (for an overview of current research, see Gleitman & Papafragou in press). Clearly, this hypothesis goes beyond the equivalence view, which is mute on the factors that give rise to the asserted correspondence between conceptual structures and lexicalization patterns. However, some empirical findings from this research are directly relevant to the issue at hand, demonstrating that, at least in some domains, the structure of conceptual knowledge does *not* vary across speakers of languages that partition the domain in different ways. The domain of motion events may serve as an example. Some languages, including English, typically encode the manner but not the path of motion in motion verbs (e.g., *stroll*, *creep*, *run*), while other languages, for example Greek, tend to encode the path rather than the manner (e.g., *anevéno*, *katevéno* for *to move up to* or *down to*, respectively). The equivalence view predicts analogous differences between English and Greek native speakers on nonverbal conceptual tasks, for example, categorizing visually presented scenes of motion events or judging their similarity. Specifically, English native speakers should pay relatively more attention to the manner than the path of the motions, while Greek native speakers should pay more attention to the path than the manner. Yet, empirical studies have found no difference between English and Greek speakers' performance on such tasks (e.g., Papafragou & Selimis 2010). Dissociations between lexicalization patterns and non-linguistic conceptual performance have also been revealed in the domain of household containers and some other domains (see Malt, Gennari & Imai 2010).

6.2. Free-concepts view.

With regard to the first problem mentioned above, an obvious alternative to the equivalence view is to acknowledge that not all concepts have a label. That is, in addition to *lexicalized concepts*, which are directly or indirectly connected with word forms, there are *free concepts*, which lack such connections. Interestingly, this view also allows one to cope with the second problem, concerning the findings from the cross-linguistic studies: People

possess a huge repertoire of concepts, which is language-independent (e.g., all people possess concepts of motion events that include manner but no path information as well as concepts that include path but no manner information). A particular subset of concepts, which is language-specific, is lexicalized. The remaining concepts are free concepts. On linguistic tasks, people use lexicalized concepts. On nonverbal tasks, they may use either free or lexicalized concepts, depending on the precise nature of the given task. By these assumptions, the free-concepts view is compatible with virtually every finding concerning the (non)correspondence between the conceptual structures of speakers of different languages, despite the fact that conceptual structures are assumed to be universal.

The free-concepts view appears awkward as it implies that people possess an enormous number of concepts, many of which differ only slightly from each other. Moreover, the proposal that people draw on a particular subset of concepts when using language is not so far from the idea that lexical meanings are knowledge structures of their own. Let us consider this view in more detail.

6.3. Non-equivalence view.

A theory postulating only one kind of knowledge structure for both lexical meanings and concepts may be an economical theory. However, whether possessing only one kind of knowledge structure for both lexical meanings and concepts is economical for the human mind is a different matter. Doubts arise when one considers the different functions that language and non-linguistic cognition have to serve. The requirements for successful communication, which is the primary function of language, differ in certain respects from what is useful for object recognition, feature prediction, action planning, and other non-linguistic cognitive tasks. Specifically, since communication is dependent on a common ground, an important constraint for lexical meanings is the shareability of the information they contain. By contrast, the information used in non-linguistic conceptual tasks need not be shared by many people. Thus, concepts may, and even should, include information deriving from a person's individual experience. In addition, lexical meanings, being based on social conventions, are probably relatively stable over time, whereas concepts should be easily malleable in order to be useful across changing life conditions. In sum, it seems reasonable to assume that concepts and lexical meanings are distinct, although related, knowledge structures (for additional arguments, see, e.g., Gleitman & Papafragou in press; Malt, Gennari & Imai 2010; Vigliocco & Vinson 2007).

One possibility is that concepts and lexical meanings are stored in distinct mental subsystems (as shown in Fig. 106.1 but are richly connected. Each concept may be linked to more than one meaning, and each meaning to more than one concept. Another possibility, which we consider more plausible, suggests itself when shifting to a finer-grained level of analysis – that is, to the constituents of concepts and lexical meanings. Conceptual and lexical-semantic knowledge may involve a common stock of atomic representations, from which they form their own more complex structures. As we have discussed in Section 3 and 4.2, there are two fundamentally different views of atomic representations in conceptual knowledge – representations of individual exemplars and representations of features (and possibly their interrelations). These two views entail different conceptions of the relationship between conceptual and lexical knowledge, which we outline in the following paragraphs.

Recall that according to the exemplar view, a concept is a set of representations of exemplars encountered in the past. In a framework that we refer to as the **common-stock-of-exemplars framework**, word meanings are assumed to comprise exemplar representations as well. However, concepts and word meanings group the representations according to different principles. A concept is a set of representations of exemplars that are similar according to non-linguistic, conceptual criteria, whereas a word meaning is the set of representations of exemplars that have been associated with this label in the past. The set of exemplar representations constituting a particular meaning may share certain exemplar representations with the set constituting a particular concept. This overlap captures the degree and the respects in which a meaning corresponds to a concept. Note that there are not two separate mental subsystems for concepts and lexical meanings, at least not in the usual sense. Rather, concepts and meanings make use of a common stock of exemplar representations. Nevertheless they are in principle independent of each other, and may for instance gradually change over time in different ways depending on particular individual experiences. Thus, this framework can account for findings concerning the (non) correspondence between lexicalization patterns and nonlinguistic conceptual organization (cf. Malt et al. 1999 for a similar framework). However, whether or not the idea that lexical meanings are sets of stored exemplars stands the test in a wider range of language processing issues has yet to be seen (see also Sec. 5.1). Up until now the exemplar approach has received little attention in psycholinguistic research.

A framework that is easier to align with common psycholinguistic assumptions is what we call the **Common-Stock-of-Features Framework**. The core idea is that both conceptual knowledge and lexical-semantic knowledge involve a common stock of microfeature representations but combine them into complex structures in a different way. This idea is instantiated in Ursino et al.'s connectionist model (Ursino, Cuppino & Magasso 2010). Instead of describing this sophisticated model in detail, let us point out some interesting aspects on the basis of a raw sketch of a model of this type (see Fig. 106.2). There are a large number of cognitive units representing elementary features which we will refer to as *microfeatures*. They are interconnected by excitatory or inhibitory connections of variable strength, with the strength of a connection reflecting the degree of statistical association or causal relationship between the respective microfeatures. Thus, conceptual knowledge is encoded in the connections among the units. A cluster of relatively strongly interconnected units can be considered a concept, but notice that there are no sharp boundaries between concepts.

The units representing microfeatures are also connected with units outside this network, for instance with units in the sensory, motor, and emotional systems (not indicated in Fig. 106.2). Furthermore, and what is most important to the present issue, with units representing the linguistic form of lexical items. For simplicity, we assume localist form representations, that is, each of these latter units represents a word form as a whole (for distributed word-form representations, see, e.g., Dilkina, McClelland & Plaut 2008). The bundle of connections between microfeatures and a particular word form make up the meaning of the respective word. Note that despite using a common stock of microfeatures, concepts and word meanings are distinct knowledge structures. Conceptual knowledge is encoded in the weights of the connections *among* the microfeatures, whereas lexical-semantic knowledge is encoded in the weights of the connections between microfeatures and forms.




Fig. 106.2: A schematic illustration of a common stock of microfeatures for conceptual and lexical-semantic knowledge. Small circles indicate units representing microfeatures, big circles indicate units representing word forms. Arrows indicate connections, with their thickness representing connection weight (Only excitatory connections having substantial weight are depicted). Dotted ellipses indicate clusters of strongly interconnected microfeatures, which may be regarded as concepts.

The cluster of microfeatures constituting a particular concept may overlap to some extent with sets of microfeatures connected to word forms. If there is a large overlap with one of these sets, then the concept may be considered a lexicalized concept. However, even if the overlap is maximal, the concept and the lexical meaning may still differ in their internal structure, since the connections that bind the microfeatures are different in the two cases. In addition, according to our previous considerations regarding specific requirements of communication, it is likely that concepts tend to be richer than the related word meanings. For example, the concept *BACHELOR* may include microfeatures representing the properties <charming> and <has no children> but possibly these microfeatures are not strongly connected with the form *bachelor*, that is, the meaning of *bachelor* may be lacking these features.

Interestingly, this framework also allows one to distinguish between homonymy (e.g., *bat*: animal / sports equipment) and polysemy (e.g., *opera*: musical drama / building). In general a word form is connected to a set of microfeatures most of which are relatively strongly connected to one another, i.e., which roughly belong to a common lump of features. In the case of homonymy, however, the word form (e.g., *bat*) is connected to different sets of microfeatures with hardly any connections between them. In the case of polysemy, the word form (e.g., *opera*) is connected to units from two or more overlapping feature lumps. The subset of features shared by the lumps captures what the different senses of the polysemous word share (e.g., has to do with music). The remaining units of the different feature lumps are barely connected to one another (cf. Rodd, Gaskell & Marslen-Wilson 2004 for a computational model of this sort).

Languages differ with regard to the composition of the sets of microfeatures that are picked out by words (e.g., a verb of motion may include or lack microfeatures representing the manner of motion). This does not a priori mean that the connections *among* the microfeatures, which make up the concepts, vary across speakers of different languages. Hence, it is possible that there is no correlation between cross-linguistic differences and performance on non-linguistic conceptual tasks. Precise predictions would require more specific assumptions about processing and learning mechanisms.

So far, we have been concerned with long-term memory structures (conceptual knowledge, lexical knowledge). Let us now briefly consider what the common-stock-of-features framework suggests with respect to working memory representations (for the distinction between long-term and working memory, see Sec. 1). Activation is fed into the microfeature network by various systems, including the sensory systems. For example, hearing a sound (e.g., barking) gives rise to the activation of units representing the features of this sound. These units in turn activate or deactivate other units, depending on whether the respective connections are excitatory or inhibitory. This complex process eventually leads to a relatively stable pattern of activated microfeatures, which includes not only microfeatures activated by the sensory input but also microfeatures strongly associated with them. Such a relatively stable pattern of activated microfeatures is what we consider a working memory representation. In our example, the working memory representation resulting from hearing the barking may include many microfeatures representing properties typical of dogs.

As to linguistic input, we assume that hearing or reading a particular word (e.g., “dog”) leads to an activation of the corresponding word-form unit. This unit in turn activates the microfeatures to which it is connected and thereby triggers a complex process in the network similar to the one described above. Again, the result is a particular pattern of activated microfeatures. This pattern is the working memory representation of the meaning of the given utterance (“dog”). Note that in this framework, word meanings are not just sets of microfeatures but rather knowledge structures that control the mapping of word forms to microfeatures (and – if we take language production into consideration – the reverse mapping as well).

It should further be noted that working memory representations derived from linguistic input do not differ in principle from ones derived from nonlinguistic input. Thus, although concepts and word meanings are clearly distinct knowledge structures, there is no analogous distinction for working memory representations in the common-stock-of-features framework. In fact, it is likely that information from various sources mix in working memory representations. Specifically, working memory representations derived from linguistic input are probably always contaminated by conceptual knowledge, as the microfeatures activated by a word form activate other microfeatures via their connections. This converges with the common belief that the linguistic meaning of a word or sentence underdetermines what listeners construe as the meaning of the utterance. Similarly, representations of nonlinguistic stimuli may be influenced by lexical knowledge, as the microfeature units activated by a nonlinguistic stimulus may feed activation to some word-form units, which in turn may feed activation back to microfeature units (cf. Papafragou & Selimis 2010 and article 107 (Landau) *Space in semantics and cognition* for on-line, transient effects of language). Finally, it should be borne in mind that conceptual and linguistic processing always takes place in a particular situation. The sensory-motor and higher cognitive systems constantly feed information into the network concerning the immediate physical environment, social situation, and current goals. Thus, working memory representations, no matter whether constructed for a nonverbal stimulus or an utterance, are always influenced by the situational context (cf. Barsalou 2009).

To summarize, the assumption that concepts are lexical meanings should be abandoned. Minimally, one needs to concede that people possess concepts that do not correspond to lexical meanings. However, it is more plausible that conceptual knowledge and lexical-semantic knowledge are distinct. This does not rule out that they are structurally and

functionally closely interwoven with one another. It is even conceivable that conceptual knowledge and lexical-semantic knowledge involve the same set of atomic representations, and only differ in how they combine these atomic representations.

7. Concluding remarks.

Research on conceptual knowledge is concerned with a wide spectrum of topics (of which we have addressed but a few). This is of little surprise considering the role of conceptual knowledge in mediating between perception, action, and higher cognitive processes. However, research on conceptual knowledge also reveals considerable heterogeneity of the fundamental views of what concepts essentially are. Do they encode rules of categorization, lists of features of category members, central-tendency information, domain-specific mini-theories, sets of exemplars encountered in the past, or bindings of modality-specific memory traces of past situations? None of these views can be dismissed. Each is supported by some empirical evidence, even though this evidence often comes from a limited type of experimental settings or paradigms. What should we conclude from this? Recent developments in research offer a promising perspective. Researchers are beginning to question some of the traditional tenets that may have hindered the development of an integrative framework for the different notions.

One traditional tenet is that the most important function of conceptual knowledge is to allow for categorization. However, in recent years an increasing number of researchers have recognized that conceptual knowledge serves many different functions besides categorization, including prediction (which we emphasized), explanation, and that different functions make use of different kinds of information (see, e.g., Barsalou 2009; Markman & Ross 2003). This calls into question another fundamental belief in research on conceptual knowledge, namely that conceptual knowledge comes in discrete packets of information, in concepts. Why should we possess discrete packets of information, accessible in an all-or-none fashion, if depending on the situation or task at hand, different pieces of these packets become relevant? Wouldn't a large network of knowledge be more plausible, from which the relevant pieces can be selected on the fly and tailored to the particular goals in the current situation? After all, even novel categories can easily be formed if necessary (see, e.g., Barsalou 1985 on *ad hoc* and goal-derived categories). Adopting this view of conceptual knowledge implies that a reframing of research in this domain is necessary. Instead of attempting to characterize the contents of concepts, the objective would be to discover what kinds of information about things is used in different situations (see Malt & Sloman 2007, Murphy & Ross 2010, and Papafragou & Selimis 2010, as examples of this approach). Research may eventually show that the different concept theories apply to different types of situations and tasks.

Intimately connected with the two aforementioned tenets is the assumption that concepts are word meanings. Indeed, a main reason why categorization is traditionally regarded a central function of conceptual knowledge is probably that concepts are taken to be word meanings. Ironically, it is exactly this assumption of a close relationship between conceptual knowledge and language that brought about the unfortunate situation that conceptual research is largely uninteresting for semantics, apart from some basic ideas (see Sec. 1). However, the situation may change as the dominance of the notion of distinct concepts gradually wanes and researchers begin to take a closer look at what pieces of information are drawn from conceptual knowledge in performing a particular task and

what the resulting working memory representation is like. Studies of this type, focusing on the conceptualization of things, situations and events in working memory are highly relevant for the question of how conceptually shaped representations are mapped onto linguistic structures and vice versa. They may therefore provide the basis for re-thinking the relationship between conceptual and lexical-semantic knowledge.

We thank Susan Olsen, Paul Portner, and Claudia Maienborn for helpful comments to an earlier version of this article.

8. References.

- Ahn, Woo-kyoung, Jesseca K. Marsh, Christian C. Luhmann & Kevin Lee 2002. Effect of theory-based feature correlations on typicality judgments. *Memory & Cognition* 30, 107–118.
- Anderson, John R. 1991. The adaptive nature of human categorization. *Psychological Review* 98, 409–429.
- Armstrong, Sharon L., Lila R. Gleitman & Henry Gleitman 1983. What some concepts might not be. *Cognition* 13, 263–308.
- Ashby, F. Gregory & W. Todd Maddox 2011. Human category learning 2.0. *Annals of the New York Academy of Sciences* 1224, 147–161.
- Barr, Robin A. & Leslie J. Caplan 1987. Category representations and their implications for category structure. *Memory & Cognition* 15, 397–418.
- Barsalou, Lawrence W. 1985. Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 11, 629–654.
- Barsalou, Lawrence W. 2009. Simulation, situated conceptualization, and prediction. *Philosophical Transactions of the Royal Society of London: Biological Sciences* 364, 1281–1289.
- Billman, Dorrit & James Knutson 1996. Unsupervised concept learning and value systematicity: A complex whole aids learning the parts. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22, 458–475.
- Close, James, Ulrike Hahn, Carl J. Hodgetts & Emmanuel M. Pothos 2010. Rules and similarity in adult concept learning. In: D. Mareschal, P. C. Quinn & S. E. G. Lea (eds.). *The Making of Human Concepts*. Oxford: Oxford University Press, 29–51.
- Collin, Charles A. & Patricia A. McMullen 2005. Subordinate-level categorization relies on high spatial frequencies to a greater degree than basic-level categorization. *Perception and Psychophysics* 67, 354–364.
- Connolly, Andrew C., Jerry A. Fodor, Lila R. Gleitman & Henry Gleitman 2007. Why stereotypes don't even make good defaults. *Cognition* 103, 1–22.
- Costello, Fintan J., & Mark T. Keane 2001. Testing two theories of conceptual combination: Alignment versus diagnosticity in the comprehension and production of combined concepts. *Journal of Experimental Psychology: Learning, Memory & Cognition* 27, 255–271.
- Davis, Tyler & Bradley C. Love 2010. Memory for category information is idealized through contrast with competing options. *Psychological Science* 21, 234–242.
- Diesendruck, Gil & Susan A. Gelman 1999. Domain differences in absolute judgments of category membership: Evidence for an essentialist account of categorization. *Psychonomic Bulletin and Review* 6, 338–346.
- Dilkina, Katia, James L. McClelland & David C. Plaut 2008. A single-system account of semantic and lexical deficits in five semantic dementia patients. *Cognitive Neuropsychology* 25, 136–164.
- Estes, Zachary 2004. Confidence and gradedness in semantic categorization: Definitely somewhat artifactual, maybe absolutely natural. *Psychonomic Bulletin and Review* 11, 1041–1047.

- Estes, Zachary & Sam Glucksberg 2000. Interactive property attribution in concept combination. *Memory & Cognition* 28, 28–34.
- Estes, Zachary, Sabrina Golonka & Lara L. Jones 2011. Thematic thinking: The apprehension and consequences of thematic relations. In: B. H. Ross (ed.). *The Psychology of Learning and Motivation*, vol. 54. Amsterdam: Elsevier, 249–294.
- Franks, Bradley 1995. Sense generation: A “quasi-classical” approach to concepts and concept combination. *Cognitive Science* 19, 441–505.
- Gagné, Christina L. & Edward J. Shoben 1997. Influence of thematic relations on the comprehension of modifier-noun combinations. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 23, 71–87.
- Gagné, Christina L. & Thomas L. Spalding 2009. Constituent integration during the processing of compound words: Does it involve the use of relational structures? *Journal of Memory and Language* 60, 20–35.
- Gelman, Susan A. 2004. Psychological essentialism in children. *Trends in Cognitive Sciences* 8, 404–409.
- Gleitman, Lila R. & Anna Papafragou in press. New perspectives on language and thought. In: K. J. Holyoak & R. G. Morrison (eds.). *The Cambridge Handbook of Thinking and Reasoning*. Cambridge: Cambridge University Press.
- Goldwater, Micah B., Arthur B. Markman & C. Hunt Stilwell 2011. The empirical case for role-governed categories. *Cognition* 118, 359–376.
- Gosselin, Frederic & Philippe G. Schyns 2001. Why do we SLIP to the basic level? Computational constraints and their implementation. *Psychological Review* 108, 735–758.
- Hampton, James A. 1982. A demonstration of intransitivity in natural categories. *Cognition* 12, 151–164.
- Hampton, James A. 1987. Inheritance of attributes in natural concept conjunctions. *Memory & Cognition* 15, 55–71.
- Hampton, James A. 1993. Prototype models of concept representation. In: I. van Mechelen et al. (eds.). *Categories and Concepts: Theoretical Views and Inductive Data Analysis*. New York: Academic Press, 67–95.
- Hampton, James A. 1997. Conceptual combination. In: K. Lamberts & D. Shanks (eds.). *Knowledge, Concepts, and Categories*. Cambridge, MA: The MIT Press, 133–159.
- Hampton, James A. 1998. Similarity-based categorization and fuzziness of natural categories. *Cognition* 65, 137–165.
- Hampton, James A. 2006. Concepts as prototypes. In: B. H. Ross (ed.). *The Psychology of Learning and Motivation*, vol. 46. San Diego, CA: Academic Press, 79–113.
- Hampton, James A. 2007. Typicality, graded membership, and vagueness. *Cognitive Science* 31, 355–384.
- Hayes, Brett K., Evan Heit & Haruka Swendsen 2010. Inductive reasoning. *Wiley Interdisciplinary Reviews: Cognitive Science* 1, 278–292.
- Johnson, Kathy E. & Carolyn B. Mervis 1997. Effects of varying levels of expertise on the basic level of categorization. *Journal of Experimental Psychology: General* 126, 248–277.
- Jolicoeur, Pierre, Mark A. Gluck & Stephen M. Kosslyn 1984. Pictures and names: Making the connection. *Cognitive Psychology* 16, 243–275.
- Jurafsky, Dan 2003. Probabilistic modeling in psycholinguistics: Linguistic comprehension and production. In: R. Bod, J. Hay & S. Jannedy (eds.). *Probabilistic Linguistics*. Cambridge, MA: The MIT Press, 39–95.
- Kalish, Charles W. 1995. Essentialism and graded membership in animal and artifact categories. *Memory & Cognition* 23, 335–353.
- Kamp, Hans & Barbara L. Partee 1995. Prototype theory and compositionality. *Cognition* 57, 129–191.

- Lazareva, Olga F. & Edward A. Wasserman 2008. Categories and concepts in animals. In: R. Menzel (ed.). *Learning and Memory: A comprehensive Reference, vol. 1: Learning Theory and Behavior*. Oxford: Elsevier, 197–226.
- Levy, Simon D. & Ross W. Gayler (eds.) 2004. *Compositional Connectionism in Cognitive Science*. Menlo Park, CA: AAAI Press.
- Love, Bradley C., Douglas L. Medin & Todd M. Gureckis 2004. SUSTAIN: A network model of category learning. *Psychological Review* 111, 309–332.
- Lynch, Elizabeth B., John D. Coley & Douglas L. Medin 2000. Tall is typical: Central tendency, ideal dimensions, and graded category structure among tree experts and novices. *Memory & Cognition* 28, 41–50.
- Maguire, Phil & Arthur W. Cater 2005. Interpreting noun-noun compounds with rare modifiers. In: K. Opwis & I.-K. Penner (eds.). *Proceedings of KogWis05. The German Cognitive Science Conference 2005*. Basel: Schwabe, 131–136.
- Maguire, Phil, Edward J. Wisniewski & Gert Storms 2010. A corpus study of semantic patterns in compounding. *Corpus Linguistics and Linguistic Theory* 6, 49–73.
- Malt, Barbara C., & Steven A. Sloman 2007. Artifact categorization: The good, the bad, and the ugly. In: E. Margolis & S. Laurence (eds.). *Creations of the Mind: Theories of Artifacts and Their Representation*. Oxford: Oxford University Press, 85–123.
- Malt, Barbara C., Silvia Gennari & Mutsumi Imai 2010. Lexicalization patterns and the world-to-word mapping. In: B.C. Malt & P. Wolff (eds.). *Words and the Mind: How Words Capture Human Experience*. Oxford: Oxford University Press, 29–57.
- Malt, Barbara C., Steven A. Sloman, Silvia Gennari, Meiyi Shi & Yuan Wang 1999. Knowing versus naming: Similarity and the linguistic categorization of artifacts. *Journal of Memory and Language* 40, 230–262.
- Margolis, Eric & Stephen Laurence (eds.) 2007. *Creations of the Mind: Theories of Artifacts and Their Representation*. Oxford: Oxford University Press.
- Markman, Arthur B. & Brian H. Ross 2003. Category use and category learning. *Psychological Bulletin* 129, 592–613.
- Martin, Alex & Alfonso Caramazza 2003. Neuropsychological and neuroimaging perspectives on conceptual knowledge: An introduction. *Cognitive Neuropsychology* 20, 195–212.
- McCloskey, Michael E. & Sam Glucksberg 1978. Natural categories: Well-defined or fuzzy sets? *Memory & Cognition* 6, 462–472.
- McNorgan, Chris, Rachel A. Kotack, Deborah C. Meehan & Ken McRae 2007. Feature-feature causal relations and statistical co-occurrences in object concepts. *Memory & Cognition* 35, 418–431.
- McRae, Ken 2004. Semantic Memory: Some insights from feature-based connectionist attractor networks. In: B. H. Ross (ed.). *The Psychology of Learning and Motivation, vol. 45*. New York: Academic Press, 41–86.
- Medin, Douglas L. & Andrew Ortony 1989. Psychological essentialism. In: S. Vosniadou & A. Ortony (eds.). *Similarity and Analogical Reasoning*. Cambridge: Cambridge University Press, 179–195.
- Medin, Douglas L. & Marguerite M. Schaffer 1978. Context theory of classification learning. *Psychological Review* 85, 207–238.
- Medin, Douglas L. & Edward J. Shoben 1988. Context and structure in conceptual combination. *Cognitive Psychology* 20, 158–190.
- Medin, Douglas L., Norbert O. Ross, Scott Atran, Douglas Cox, John Coley, Julia B. Profitt & Sergey Blok 2006. Folkbiology of freshwater fish. *Cognition* 99, 237–273.
- Mervis, Carolyn B. & Eleanor Rosch 1981. Categorization of natural objects. *Annual Review of Psychology* 32, 89–115.
- Minda, John Paul & J. David Smith 2002. Comparing prototype-based and exemplar-based accounts of category learning and attentional allocation. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 28, 275–292.

- Minda, John Paul & J. David Smith 2011. Prototype models of categorization: Basic formulation, predictions, and limitations. In: E. M. Pothos & A. J. Wills (eds.). *Formal Approaches in Categorization*. Cambridge: Cambridge University Press, 40–64.
- Moss, Helen E., Lorraine K. Tyler & Kirsten I. Taylor 2007. Conceptual structure. In: M. G. Gaskell (ed.). *The Oxford Handbook of Psycholinguistics*. Oxford: Oxford University Press, 217–234.
- Murphy, Gregory L. 1988. Comprehending complex concepts. *Cognitive Science* 12, 529–562.
- Murphy, Gregory L. 1990. Noun phrase interpretation and conceptual combination. *Journal of Memory and Language* 29, 259–288.
- Murphy, Gregory L. 2002. *The big book of concepts*. Cambridge, MA: The MIT Press.
- Murphy, Gregory L., & Mary E. Lassaline 1997. Hierarchical structure in concepts and the basic level of categorization. In: K. Lamberts & D. Shanks (eds.). *Knowledge, Concepts, and Categories*. Cambridge, MA: The MIT Press, 93–131.
- Murphy, Gregory L. & Douglas L. Medin 1985. The role of theories in conceptual coherence. *Psychological Review* 92, 289–316.
- Murphy, Gregory L. & Brian H. Ross 2010. Category vs. object knowledge in category-based induction. *Journal of Memory and Language* 63, 1–17.
- Nosofsky, Robert M. 1986. Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General* 115, 39–57.
- Nosofsky, Robert M. 2000. Exemplar representation without generalization? Comment on Smith and Minda's (2000) "Thirty categorization results in search of a model". *Journal of Experimental Psychology: Learning, Memory, and Cognition* 26, 1735–1743.
- Onishi, Kristine H., Gregory L. Murphy & Kathryn Bock 2008. Prototypicality in sentence production. *Cognitive Psychology* 56, 103–141.
- Osherson, Daniel N. & Edward E. Smith 1981. On the adequacy of prototype theory as a theory of concepts. *Cognition* 9, 35–58.
- Osherson, Daniel N. & Edward E. Smith 1997. On typicality and vagueness. *Cognition* 64, 189–206.
- Papafragou, Anna 2005. Relations between language and thought: Individuation and the count/mass distinction. In: H. Cohen & C. Lefebvre (eds.). *Handbook of Categorization in Cognitive Science*. Amsterdam: Elsevier, 255–275.
- Papafragou, Anna & Stathis Selimis 2010. Event categorisation and language: A cross-linguistic study of motion. *Language and Cognitive Processes* 25, 224–260.
- Peeters, Bert (ed.). 2000. *The Lexicon-Encyclopedia Interface*. Amsterdam: Elsevier.
- Posner, Michael I. & Steven W. Keele 1968. On the genesis of abstract ideas. *Journal of Experimental Psychology* 77, 353–363.
- Pothos, Emmanuel M. & Andy J. Wills (eds.) 2011. *Formal Approaches in Categorization*. Cambridge: Cambridge University Press.
- Rakison, David H. & Yevdokiya Yermolayeva 2010. Infant categorization. *Wiley Interdisciplinary Reviews: Cognitive Science* 1, 894–905.
- Rehder, Bob 2010. Causal-based classification: A review. In: B. H. Ross (ed.). *The Psychology of Learning and Motivation*, vol. 52. Amsterdam: Elsevier, 39–116.
- Rips, Lance J. 1995. The current status of research on concept combination. *Mind & Language* 10, 72–104.
- Rodd, Jennifer M., M. Gareth Gaskell & William D. Marslen-Wilson 2004. Modelling the effects of semantic ambiguity in word recognition. *Cognitive Science* 28, 89–104.
- Rogers, Timothy T. & Karalyn Patterson 2007. Object categorization: Reversals and explanations of the basic-level advantage. *Journal of Experimental Psychology: General* 136, 451–469.
- Rosch, Eleanor 1978. Principles of categorization. In: E. Rosch & B. B. Lloyd (eds.). *Cognition and Categorization*. Hillsdale, NJ: Erlbaum, 27–48.
- Rosch, Eleanor, & Carolyn B. Mervis 1975. Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology* 7, 573–605.

- Rosch, Eleanor, Carolyn B. Mervis, Wayne D. Gray, David M. Johnson & Penny Boyes-Braem 1976. Basic objects in natural categories. *Cognitive Psychology* 8, 382–439.
- Roth, Emilie M. & Carolyn B. Mervis 1983. Fuzzy set theory and class inclusion relations in semantic categories. *Journal of Verbal Learning and Verbal Behavior* 22, 509–525.
- Schyns, Philippe G., Robert L. Goldstone & Jean-Pierre Thibaut 1998. The development of features in object concepts. *Behavioral and Brain Sciences* 21, 1–54.
- Shin, Hyun Jung & Robert M. Nosofsky 1992. Similarity-scaling studies of dot-pattern classification and recognition. *Journal of Experimental Psychology: General* 121, 278–304.
- Sloman, Steven A. 1998. Categorical inference is not a tree: The myth of inheritance hierarchies. *Cognitive Psychology* 35, 1–33.
- Smith, Edward E. & Douglas L. Medin 1981. *Categories and Concepts*. Cambridge, MA: Harvard University Press.
- Smith, Edward E., Daniel N. Osherson, Lance J. Rips & Margaret Keane 1988. Combining prototypes: A selective modification model. *Cognitive Science* 12, 485–527.
- Smith, J. David & John Paul Minda 2000. Thirty categorization results in search of a model. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 26, 3–27.
- Smith, Linda B. & Eliana Colunga in press. Developing categories and concepts. In: M. Spivey, M. Joanisse & K. McRae (eds.). *The Cambridge Handbook of Psycholinguistics*. Cambridge: Cambridge University Press.
- Snedeker, Jesse & Lila R. Gleitman 2004. Why it is hard to label our concepts. In: D. G. Hall & S. R. Waxman (eds.). *Weaving a Lexicon*. Cambridge, MA: The MIT Press, 257–293.
- Storms, Gert 2004. Exemplar models in the study of natural language concepts. In: B. H. Ross (ed.). *The Psychology of Learning and Motivation, vol. 45*. New York: Academic Press, 1–39.
- Storms, Gert, & Edward J. Wisniewski 2005. Does the order of head noun and modifier explain response times in conceptual combination? *Memory & Cognition* 33, 852–861.
- Storms, Gert, Paul de Boeck, Iven van Mechelen & Dirk Geeraerts 1993. Dominance and noncommutativity effects in concept conjunctions: Extensional or intensional basis? *Memory & Cognition* 21, 752–762.
- Swinney, David, Tracy Love, Matthew Walenski & Edward E. Smith 2007. Conceptual combination during sentence comprehension. *Psychological Science* 18, 397–400.
- Tanaka, James W. & Marjorie Taylor 1991. Object categories and expertise: Is the basic level in the eye of the beholder? *Cognitive Psychology* 23, 457–482.
- Ursino, Mauro, Cristiano Cuppino & Elisa Magosso 2010. A computational model of the lexical-semantic system, based on a grounded cognition approach. *Frontiers in Psychology* 1:221. doi: 10.3389/fpsyg.2010.00221.
- Vanpaemel, Wolf & Gert Storms 2008. In search of abstraction: The varying abstraction model of categorization. *Psychonomic Bulletin and Review* 15, 732–749.
- Vigliocco, Gabriella & David P. Vinson 2007. Semantic representation. In: M. G. Gaskell (ed.). *The Oxford Handbook of Psycholinguistics*. Oxford: Oxford University Press, 195–215.
- Wilkenfeld, Merryll J. & Thomas B. Ward 2001. Similarity and emergence in conceptual combination. *Journal of Memory and Language* 45, 21–38.
- Wisniewski, Edward J. 1997. When concepts combine. *Psychonomic Bulletin and Review* 4, 167–183.
- Wisniewski, Edward J. 2009. On using count nouns, mass nouns, and *pluralia tantum*: What counts? In: F. J. Pelletier (ed.). *Kinds, Things, and Stuff: Mass Terms and Generics*. Oxford: Oxford University Press, 166–191.
- Wisniewski, Edward J., Mutsumi Imai & Lyman Casey 1996. On the equivalence of superordinate concepts. *Cognition* 60, 269–298.
- Wu, Ling-Ling & Lawrence W. Barsalou 2009. Perceptual simulation in conceptual combination: Evidence from property generation. *Acta Psychologica* 132, 173–189.

Stephanie Kelter, Berlin (Germany)
Barbara Kaup, Tübingen (Germany)