

Simple Heuristics for Concept Combination

Edouard Machery & Lisa G. Lederer

For M. Werning, W. Hinzen, and E. Machery (Eds.)

In psychology, concepts are typically characterized as those bodies of knowledge that are used by default in the psychological processes underlying numerous higher cognitive competences, such as categorization, induction, language understanding, analogy-making, and so on (Machery, In Press). These cognitive competences include concept combination, the capacity to produce new concepts out of pre-existing concepts—for instance, the capacity to produce the concept of a Harvard graduate who is a carpenter out of the concepts of a Harvard graduate and of a carpenter (Kunda, Miller, & Clare, 1990). The outcome of this process is said to be a “complex concept,” while the concepts permanently stored in long-term memory are said to be “simple concepts.”¹

Psychologists interested in concepts have studied at length the psychological phenomena associated with concept combination in order to better understand the cognitive processes underlying concept combination (for an overview, see Murphy, 2002, Chapter 12). However, psychological research on concept combination has failed to yield any consensus on the nature of these processes. In this article, we will critically review the main theories of concept combination that have been developed since the 1980s. This review will set the stage for some new hypotheses about the cognitive processes underlying concept combination, which emphasize their diversity and their ecological rationality.

Here is how we will proceed. In the first three sections, we consider three important models of concept combination—Smith and colleagues’ Selective Modification Model (Section 1), Hampton’s Composite Prototype Model (Section 2), and Costello and

¹ It is noteworthy that the distinction between simple and complex concepts does not exactly map onto the distinction between simple and complex expressions. A complex expression, such as “soccer moms,” can express a concept permanently stored in long-term memory, thus a simple concept, according to the terminology used in this article.

Keane's C^3 model (Section 3).² In Section 4, we contrast these models with the insights about cognitive processing that emerge from Gigerenzer and colleagues' "Fast and Frugal Heuristics" research program. In Section 5, we apply these insights to concept combination to propose several new hypotheses about the processes underlying the creation of complex concepts.

1. The Selective Modification Model of Concept Combination

1.1 The Selective Modification Model

Smith and colleagues' famous model of concept combination combines a model for producing complex concepts out of simple concepts with a prototype model of concept representation and a metric for computing the typicality of objects with respect to those concepts (Smith, Osherson, Rips, & Keane, 1988). We consider the first two components in turn.

Prototypes, roughly, are bodies of statistical knowledge about classes of physical objects, types of events, or substances. Smith and colleagues' version of the prototype theory of concepts distinguishes attributes from values. In substance, attributes are kinds of property, while values are properties. *Color* is an attribute, while *red* and *blue* are values. Smith and colleagues propose that for each attribute it represents, a concept stores some knowledge about the distribution of properties among the members of its extension (Figure 1). For example, the prototype of apple is assumed to store some knowledge about how often apples are red, how often apples are blue, how often apples are green,

² Because some of these models are further described in other chapters of the *Oxford Handbook of Compositionality*, we only focus on their most important properties. For further detail, the reader should consult Prinz's chapter XX, Hampton and Jönsson's chapter XX, and Connolly, Armstrong, and Gleitman's chapter XX. For the sake of space, Wisniewski's Dual-Process Model is not described here, but many of the points made in this chapter also apply to it (see Wisniewski's chapter XX).

and so on.³ Each value has a certain number of “votes” reflecting its salience, which in turn is supposed to reflect its subjective frequency. If someone assumes that 80% of apples are red, then the value *red* for the attribute *color* should have 80% of the sum of the votes for the attribute *color*. Attributes vary in their diagnosticity, which is defined as “a measure of how useful the attribute is in discriminating instances of the concept from instances of contrasting concepts” (1988, 487).

Put Figure 1 about here.

The model of concept combination proposed by Smith and colleagues applies only to modifier-head complex concepts. In the case of a modifier-head complex concept, a concept such as BLUE modifies a head concept, such as APPLE, forming the complex concept BLUE APPLE.⁴ Smith and colleagues propose that concept combination consists in modifying the representation of the head concept according to the following procedure. Modifications are limited to the attribute that corresponds to the property expressed by the modifier concept. The remainder of the representation remains unchanged (except for the fact that the relative diagnosticity of the other attributes decreases—see below). For instance, for the complex concept BLUE APPLE, the modifications of the concept APPLE are limited to the attribute *color*. The concept APPLE represents the various colors of apples, each color’s votes corresponding to its subjective frequency (see Figure 1). By contrast, the concept BLUE APPLE represents all the members of its extension (all blue apples) as being blue. That is, all the values for the attribute *color* in the concept APPLE are replaced with a single value—viz., *blue*. The number of votes for the value *blue* in the complex concept BLUE APPLE is then equal to the combined number of votes for all the color values in the simple concept APPLE (see Figure 2). The diagnosticity of the modified attribute *color* is also increased.

³ Smith and colleagues do not specify whether the represented values include only those values that have actually been encountered (e.g., the value *red* for the apples) or all of the values that are conceivable (e.g. the value *blue* for the apples).

⁴ We use small caps to name concepts and italics to name properties. We also use the symbol “A^B” to name the complex concept produced by combination of the concepts A and B.

Put Figure 2 about here.

Above and beyond the details and mechanics of the Selective Modification model, it is worthwhile highlighting the key insights about concept combination it incorporates. The central idea is that people assume that members of the extension of a modifier-head complex concept are just like the members of the extension of the head concept but for one respect (corresponding to the property expressed by the modifier). Thus, the model assumes that people assume blue apples to be just like apples, except that they are all blue (while some apples are red, green, etc.). If people believe that most apples are sour, they should believe, according to the Selective Modification model, that most blue apples are sour. The second insight built into this model is that when we reason with a complex concept (including when we categorize according to a complex concept), the attribute corresponding to the modifier is more important than when we reason with the original head concept. For instance, when we decide whether to categorize an object as a blue apple, the color of this object is more important than when we merely decide whether to categorize it as an apple. The greater diagnosticity of the modified attribute captures this insight.

1.2 Empirical Adequacy

The Selective Modification model explains successfully various findings about objects' perceived typicality with respect to the classes denoted by simple and complex concepts (Smith et al., 1988). According to this model, an object's typicality with respect to the extension of a concept is an increasing function of the properties that are both possessed by the object and represented by the concept (e.g., the redness of a red apple), and a decreasing function of the properties that are represented by the concept (e.g., green for the attribute *color*), but not possessed by the object (the red apple), weighted by the vote of each value (e.g., *red* and *green*) and the diagnosticity of each attribute (e.g., *color*). Thus, the model predicts (1) that a brown apple will be judged to be less typical of green apples than of apples and (2) that a green apple will be judged to more typical of green apples than of apples. People's judgments of typicality confirm these predictions.

1.3 Empirical Objections Against the Selective Modification Model

The empirical shortcomings of the Selective Modification model have been extensively studied. Medin and Shoben (1988) and Murphy (1988) have provided evidence against the hypothesis that an instance of a modifier-head complex concept is taken to be identical to an instance of the head concept in every respect but one (corresponding to the property expressed by the modifier). Particularly, Medin and Shoben have shown that combining a head concept, e.g., SPOON, with a modifier concept, e.g., WOODEN, does not simply affect the attribute corresponding to the property expressed by the modifier concept (*viz. material*), but other attributes as well. They asked participants to evaluate the typicality of some objects with respect to the classes denoted by both a head concept class and a complex concept formed out of this head concept. For example, participants were asked to evaluate the typicality of small and large spoons with respect to both spoons and wooden spoons. The Selective Modification model predicts that the representation of wooden spoons should be identical to the representation of spoons, except for the material attribute; the *size* attribute should not be affected. However, while people found large spoons, compared to small spoons, to be more typical of wooden spoons, they found small spoons to be more typical of spoons in general (for other findings, see Medin & Shoben, 1988, Murphy, 1988, and Murphy, 1990).

1.4 Theoretical Objections Against the Selective Modification Model

It is clear that the scope of the selective modification model is limited. It bears only on complex concepts that combine a modifier concept with a head concept. Nothing is said about other types of complex concepts, such as the complex concepts expressed by noun-noun complex expressions such as “apartment dog” or “house bird” (Murphy, 1988), or the modifier + head complex concepts, in which the modifier concept, unlike BLUE, does not specify the same attribute when paired with all head concepts. For example, like all concepts expressed by non-predicating adjectives, the concept expressed by “corporate” modifies differently the head concepts LAWYER and BUILDING (Levi, 1978; Murphy, 1988, 1990). We can also combine head concepts with modifier concepts that express properties for which the head concept does not have any relevant attribute. Consider for instance the complex concept expressed by “summer smile”: it is unlikely that SMILE has any attribute with values corresponding to the four seasons.

Furthermore, some aspects of the Selective Modification Model are left unspecified. When we produce a complex concept such as BLUE APPLE, we are supposed to increase the diagnosticity of the attribute color. However, Smith and colleagues do not specify how diagnosticity is to be increased, even though the empirical adequacy of the model depends on how exactly it is increased. This increase has to be sufficiently large so that a strange blue apple (say, a squarish blue apple) is a better example of a blue apple than a typical green apple is. Simultaneously, however, it has to be sufficiently small so that a blue pear is a worse example of a blue apple than of a pear.

2. Hampton's Composite Prototype Model of Concept Combination

2.1 The Composite Prototype Model

Hampton's model of concept combination includes, along with a qualitative description of the process resulting in complex concepts, a model of concepts, a model of categorization, and a model of typicality judgment (Hampton, 1982, 1987, 1988, 1996, 1997a, b; Storms, de Boeck, Hampton, & van Mechelen, 1999).⁵ We will focus mostly on the first two components of this model. Just like Smith and colleagues, Hampton assumes that concepts are prototypes. In his prototype model of concepts, a concept consists simply of a list of weighted properties (Figure 3). Like the values of Smith and colleagues' model, properties are weighted more heavily as their subjective frequency among members of their concept category becomes greater. A property's weight is thus a measure of its typicality.

Put Figure 3 about here.

For Hampton, concept combination is a constrained process of the inheritance of typical properties from the combined concepts (say, GRANDMOTHER and SPY) to the complex concept (GRANDMOTHER SPY). That is, first, when they produce a complex concept, people assume by default that the properties they judge to be typical of the classes represented by the combined concepts are also typical of the class represented by the complex concept. Thus, the properties that are judged to be typical of grandmothers

⁵ This model has inspired Prinz's (2002) and Machery's (2005) models of concept combination.

and of spies are by default judged to be typical of grandmother spies. Hampton calls this characteristic of concept combination “attribute inheritance.” Attribute inheritance is related to Smith and colleagues’ idea that an instance of a modifier + head concept is similar to an instance of a head concept in all respects but one (corresponding to the dimension of the head concept modified by the modifier concept). Both models assume that by default, we view a blue apple as having the typical properties of apples.

Second, in Hampton’s model, property inheritance obeys some constraints. First, in some combinations, the typical properties of one of the combined concepts are more likely to be inherited by the complex concept than the typical properties of the other combined concept. In such cases, an instance of a complex concept $A \wedge B$ is judged to be an instance of one of the combined concepts (say, A) that has a relatively small number of properties of the instances of the concept B. Hampton calls this “the concept dominance effect.” Second, not every property typical of a combined concept is inherited by the complex concept. When a property typical of the instances of one of the combined concepts is judged to be incompatible with membership in the class represented by the other one, this property is not inherited by the complex concept. Third, properties judged to be necessary for membership in the class represented by one of the combined concepts are always inherited by the complex concept.

In addition, Hampton also proposes that complex concepts represent what he calls “emergent properties”, viz. properties represented by the complex concept that are not derived from either of the combined concepts. For instance, Harvard graduates who are carpenters might be judged to be idealistic, even though neither Harvard graduates nor carpenters are judged to be idealistic (Kunda et al., 1990). Such emergence derives from two sources. In some cases, people who produce a complex concept are already familiar with some members of its extension and assume that most instances of the complex concept possess the same properties as these members (a process called “extensional feedback” by Hampton). For example, when they produce the complex concept PET FISH, people might rely on their memories of specific pet fish. In other cases, people might reason about why something would be an instance of the complex concept. Thus, people might reason about why a Harvard graduate might become a carpenter by hypothesizing that she was idealistic (Kunda et al., 1990; Hastie, Schroeder & Weber, 1990).

2.2 Empirical Adequacy

This model successfully explains a large number of empirical findings (for further detail, see Hampton and Jönsson's chapter in this volume). Hampton (1987) asked participants to list the properties that were typical of the instances of two concepts (A and B) as well as the properties that were typical of the instances of the resulting complex concept. He then asked other participants to evaluate how important each of the listed properties was for categorizing an object as a member of the classes denoted by A, B, and $A \wedge B$. As predicted by the Composite Modification Model, he found that the importance of a given property with respect to $A \wedge B$, according to participants' evaluations, was a positively increasing function of its importance with respect to A and to B.

This model also explains various puzzling characteristics of concept combination. Overextension might be the most striking characteristic (Hampton, 1982, 1988, 1996): people sometimes judge that an object that is not an instance of a concept A is an instance of the complex concept $A \wedge B$. For instance, people tend to judge that blackboards are not pieces of furniture, while judging that they are pieces of school furniture. This finding follows from the Composite Prototype Model on the assumption that a blackboard is not similar enough to the prototype of a piece of furniture for people to decide that a blackboard is a piece of furniture, while being similar enough to the prototype of a piece of school furniture for people to decide that it is a piece of school furniture.

2.3 Empirical Objections Against the Composite Prototype Model

Connolly, Fodor, Gleitman, and Gleitman (Connolly, Fodor, Gleitman, & Gleitman, 2007; see also Connolly et al.'s chapter in this volume) have recently challenged what they see as a central idea of Hampton's and Smith and colleagues' models—viz. the hypothesis that people assume that properties typical of the instances of concept A are equally typical of the instances of the complex concept $A \wedge B$. If Hampton's and Smith and colleagues' models were right, Connolly and colleagues argue, then adding a modifier to a noun (or a noun phrase) should not affect the judged typicality of any properties of its extension that have not been changed explicitly. As they show, however, people judge instances of nouns alone (e.g., “ducks”) to be more likely than instances of modifier +

noun combinations (e.g., “baby ducks”) to possess certain typical properties (e.g., *having webbed feet*). People also judge instances of modifier + noun combinations (e.g., “baby ducks”) to be more likely than instances of double modifier + noun combinations (e.g., “Peruvian baby ducks”) to possess these properties.

Importantly, however, Jönsson and Hampton (2008) argue that Connolly and colleagues’ study does not provide evidence against either Hampton’s or Smith and colleagues’ models. As we saw, Smith and colleagues specify that the weight of the attribute that is modified by the modifier concept is increased. This shift in weighting entails that all other attributes, whether or not they are related to the modifier, are relatively *less* weighted. Thus, creating the complex concept BABY DUCK results both in the *age* attribute in the prototype DUCK being more heavily weighted and in the *foot type* attribute being less heavily weighted. The Selective Modification model is thus compatible with Connolly and colleagues’ finding that people judge baby ducks less likely than ducks to have webbed feet. Similarly, in the Composite Prototype Model, the weight of a property in a complex concept is an increasing function of its weight in *each* combined concept. When a head concept is modified, the weights of all its properties are impacted by the corresponding weights in the modifier concept. In Connolly and colleagues’ experiment, because instances of BABY only rarely have webbed feet, participants would expect baby ducks to be less likely than ducks to have webbed feet.

3. Costello and Keane’s C³ Model of Concept Combination

3.1 The C³ Model

Costello and Keane (2000, 2005) propose a computational model and an algorithmic model of concept combination (for details about the algorithmic model, see Costello & Keane, 2000). They focus on the conceptual combination involved in interpreting novel noun-noun compounds, such as “street knife.” At the computational level of description, concept combination consists of the optimal satisfaction of three constraints: diagnosticity, plausibility, and informativeness. According to the diagnosticity constraint, the complex concept should be better expressed by the uttered noun-noun compound than by other noun-noun compounds; the more the complex concept represents its instances as having some properties that are diagnostic of the combined concepts, the more this

constraint is met. According to the plausibility constraint, the complex concept expressed by a noun-noun compound should refer to objects that are plausible in light of previous knowledge. According to the informativeness constraint, both nouns are necessary and sufficient to convey the information the speaker intends to convey. Costello and Keane argue that these three constraints are derived from the pragmatic rules governing successful communication. To satisfy them, concept combination can appeal to any body of knowledge stored in long-term memory. They propose that “the constraint theory proposes that the combination process has direct access to the full contents of memory” (Costello & Keane, 2000, 305).

3.2 Empirical Adequacy

Costello and Keane (2000) have shown that C^3 is consistent with several characteristics of the interpretation of novel noun-noun compounds, such as the existence of several forms of interpretation, the polysemy of these compounds, and the richness of the interpretations people propose. One form of interpretation of noun-noun compounds that Costello and Keane (2001) focus on is called “property interpretation” (Wisniewski, 1996). In property interpretation, the listener supposes that the first noun in the noun-noun compound (e.g., “whale” in “whale boat”) is used to ascribe a property to the referents of the second noun; for instance, a whale boat might be a very big boat. As C^3 predicts, the properties ascribed to the referents of the second noun are typically highly diagnostic of the concept expressed by the first noun.

3.3 Theoretical Objections

Although this is not the place to criticize at length C^3 , we note that Costello and Keane focus on a single environment in their account of concept combination, viz. linguistic communication, and they propose a single computation-heavy process for the task of combining concepts. We will come back to these two points in the next section.

4. Fast and Frugal Heuristics and Concept Combination

4.1 Fast and Frugal Heuristics

Gigerenzer and colleagues have developed an attractive picture of the cognitive processes underlying various cognitive competences, such as choice under uncertainty or categorization (e.g., Gigerenzer, Todd, & the ABC Research Group, 1999). According to Gigerenzer and colleagues, these cognitive processes are simple, fast, and frugal. Their speed and their simplicity result from their frugality: they do not take into account all the cues that could be relevant to make a judgment or a decision. Rather, they typically take into account a few cues and sometimes a single one (in which case they are called “one-reason decision rules”). As an example of a simple and frugal process, consider the recognition heuristic (Gigerenzer & Goldstein, 1996; Goldstein & Gigerenzer, 2002). It is a decision rule that applies to choices between two (or more) options. For choices among two options, the heuristic can be described as follows:

Recognition heuristic: If one of two objects is recognized and the other is not, then infer that the recognized object has the higher value with respect to the criterion. (Goldstein & Gigerenzer, 2002, 76)

Suppose that a European has to decide whether San Antonio or San Diego has the larger population. San Antonio and San Diego are the two options and the city size is the criterion. Suppose also that she only recognizes San Diego. If she were to apply the recognition heuristic, she would conclude correctly that San Diego is bigger than San Antonio. The recognition heuristic is a simple and frugal one-reason decision rule, since it takes into account a single cue—whether or not the options (e.g., San Diego or San Antonio) are recognized.

A second working hypothesis is that these rules are paired with specific environments in which they perform as well, and sometimes better, than more complex rules. These rules are said to be ecologically rational in these environments. Consider again the recognition heuristic. It works well when recognition is correlated with the criterion (viz. size in the example above). Because the larger a city’s size, the more recognizable it is to foreigners, the recognition heuristic is ecologically rational for Germans, but not for Americans.

A third working hypothesis is that we possess numerous rules for making judgments, for making choices, and for drawing inferences. For instance, proponents of the fast and frugal heuristics research program propose that besides the recognition

heuristic, there are other heuristics that can be used to decide which of two (or more) options has the highest value with respect to a criterion, for instance, which of two cities has the larger population. Take-the-Best is one of them. Take-the-Best is supposed to be used to choose between two options that are ranked with respect to a criterion, on the basis of cues that are more or less valid. The validity of a cue is defined as the relative frequency with which options with a positive cue have a higher value with respect to the criterion than options with a negative cue. Take-the-Best is a one-reason decision rule. We look for the most valid cue that discriminates between the two options and we decide on the basis of this cue, thereby neglecting less valid cues. For example, a person using Take-the-Best may decide whether Bonn or Dresden has the larger population based on the cue of having a soccer team in the Bundesliga. The validity of this cue is defined as the relative frequency with which cities with a soccer team in the Bundesliga are larger than cities without a soccer team in the Bundesliga. Its validity is imperfect, because it happens that the smaller of two cities has a soccer team in the Bundesliga while the larger does not. If having a soccer team is the most valid cue, and if Bonn has a soccer team while Dresden does not, a person using Take-the-Best would decide that Bonn has the larger population. If both Bonn and Dresden have both a team in the Bundesliga or if neither does, the person would move to the next cue. As Gigerenzer and colleagues put it, our cognitive processes form a toolbox of simple heuristics. They propose that for a given task, we can use several heuristics. If we do not recognize one of the options, we might use the recognition heuristic. If we do recognize both options, we might use Take-the-Best or other heuristics in the toolbox.

Proponents of the Fast and Frugal Heuristics research program have not studied the process of concept combination; so, it is unknown whether some processes of concept combination can be understood as fast and frugal heuristics. In Section 5, we will consider various hypotheses about concept combination inspired by this research program. What we want to do in the remainder of this section is to examine critically the past research on concept combination in light of the ideas summarized above. We will argue that past research has systematically overlooked the insights about cognitive processing that inspire the Fast and Frugal Heuristics research program.

4.2 A Toolbox of Ecologically Rational Heuristics

Psychologists working on concept combination sometimes recognize that there are several distinct processes of concept combination, each of which is meant to produce a distinct kind of complex concept. As we saw in Section 1, Smith and colleagues' Selective Modification Model is meant to apply only to modifier + head complex concepts, such as BLUE APPLE, and they suggest that other processes might be involved in producing other types of complex concepts. Hampton (1997) distinguishes two processes—a process described by the Composite Prototype model (see Section 2) and a process involved in the interpretation of novel noun-noun compounds, such as “trumpet olive” (see also Wisniewski, 1996, 1997).

However, psychologists' acknowledgment that there are several processes of concept combination does not go far enough. To see why, we need to remember that complex concepts are always produced *for a purpose* (see Downing, 1977; Wisniewski, 1997; Costello & Keane, 2000). Complex concepts are used to categorize, they are used to make inductive inferences, they are used to draw analogies, and so on. In fact, complex concepts are probably used in most tasks whose solution involves retrieving concepts from long-term memory. We now consider in more detail some of these contexts.

In a communicational context, complex concepts are often produced to understand the assumptions and inferences made by a speaker or writer about the instances of the complex concept (see Costello & Keane, 2000, 2005 on this point). Consider, for instance, the following extract from James Ellroy's novel *The Black Dahlia* (1987, 218):

I dumped my grip on the dresser, and as a precaution on the way out, yanked two hairs from my head and spit-glued them across the door-doorjamb juncture. If the fascisti prowled the pad, I would know.

To make sense of this passage, the reader has to assume that spit-gluing (viz. gluing with spit) has by and large the same effects as gluing and thus that the hairs that are spit-glued will remain spit-glued except if the door is open (as we saw, this idea is built into the Selective Modification Model, the Composite Prototype Model, and the C³ Model).

Concept combination also happens for merely referential purposes. For instance, a hearer might have to identify the referent of a complex linguistic compound. Nunberg's

(1979) famous “ham sandwich” example illustrates this situation: a waitress might have to identify the referent of “ham sandwich” when she is told, “The ham sandwich at the corner table wants another root beer.” In this case, in contrast to the situation described earlier, the point of concept combination is not to understand the assumptions and inferences made by a speaker or writer about the instances of a complex concept; the point for the cognizer is simply to identify an object in the environment.

Complex concepts are also produced for non-linguistic purposes. One might produce a complex concept for the purpose of categorization. For instance, one might go to the market to buy some Chinese cabbage called for in a recipe. If one does not already have a concept of Chinese cabbage, then one needs to produce one in order to be able to select the correct vegetable. The concept combination process might vary with the particular selection task. It is one thing to decide whether to classify a single object as a Chinese cabbage (as would, e.g., a person deciding whether to buy a particular vegetable that might or might not be a Chinese cabbage); it is another to decide whether to select any of several vegetables; it is yet another to decide which of several objects to select when one knows that one of the objects must be selected. Time constraints also matter. It is one thing to make a selection without any time constraint and another to have to decide quickly whether an object belongs to a given class.

The idea that concepts are combined for various specific purposes might seem trivial, but it is important to keep it in mind when one attempts to understand concept combination. We hypothesize that the cognitive processes underlying concept combination also vary, depending on the reason why concepts are combined, and not only depending on the type of the complex concepts that are produced, as many psychologists have assumed. That is, concepts might be combined differently when the complex concept is used to understand the assumptions or inferences made by a writer, when it is merely used to identify an object in one’s environment, when it is used to categorize something in a category one is not familiar with, and so on. This hypothesis follows from Gigerenzer and colleagues’ idea that a cognitive process “fits” its environment, in the sense that its structure results in satisfying outcomes in this environment (but might fail to yield any satisfying outcomes in other environments). We

hypothesize that for the processes underlying concept combination, the relevant environments are defined by the goals for which complex concepts are created.

4.3 Simple Processing vs. Isotropic Processing

Most models of concept combination assume that people can, and often do, access their whole background knowledge during concept combination (a property called “isotropy”).⁶ For example, in Hampton’s Composite Prototype Model (Section 2), emergence is explained by people accessing their knowledge about the world (extensional feedback) and engaging in causal reasoning, and no limit is put on the knowledge that is ultimately accessed.⁷ Thus, when people judge that Harvard graduates who are carpenters are idealistic, they are supposed to appeal to causal generalizations linking character traits and behavior.

The hypothesis that concept combination is underwritten by simple heuristics suggests, however, that having access to our total background knowledge during concept combination is the exception, rather than the rule. For many purposes, it might be sufficient to have access only to a limited amount of knowledge. For instance, Gagné and Shoben (Gagné & Shoben, 1997) have shown that when people interpret novel noun-noun compounds, such as “mountain magazine” and “mountain stream,” they are quicker when the relation between the two nouns is a relation that is typical for the modifier noun; because “mountain” is usually used with the thematic relation *in* rather than with the thematic relation *about*, people judge “mountain stream” to be sensible more quickly

⁶ Not all the models considered above assume that the processes of concept combination are isotropic. The Selective Modification model does not require people to access their background knowledge to produce a complex concept.

⁷ In the Composite Prototype Model, concept combination involves two stages: during the first stage, the processes of concept combination have a restricted access to the knowledge stored in long term memory, since only the information stored in prototypes is accessible; during the second stage, however, all the knowledge in long-term memory is accessible. By contrast, in C^3 , the process of concept combination accesses simultaneously the whole knowledge base in long-term memory.

than “mountain magazine.” A plausible interpretation of this finding is that people preferentially access a limited body of knowledge about the referents—including some knowledge about the typical relations these referents stand in—when they interpret a novel noun-noun compound. Only when these typical relations do not lead to a plausible interpretation are additional semantic resources considered.

Furthermore, cognitive processes that do not depend on accessing any background knowledge, such as the process described by Smith and colleagues’ Selective Modification Model, might be more common than is suggested by the general rejection of this model in the literature on concept combination (see also Gagné and Shoben’s (1997) CARIN model).

4.4 Simple Processing vs. the Production of “Comprehensive” Complex Concepts

Another respect in which past research on concept combination overlooks the hypothesis that cognitive processes can be both efficient and simple is the assumption that the processes of concept combination produce “comprehensive” complex concepts. By “comprehensive concepts,” we mean concepts that represent a large amount of information about their extension: they describe their extension in detail. Thus, a complete concept of a Harvard graduate who is a carpenter would store a large amount of information about the supposed properties of Harvard graduates who are carpenters. Comprehensive complex concepts stand in contrast with “sketchy” complex concepts, which only represent the information about their extension that is needed to fulfill the purpose of their production. Limited knowledge access, as described above, leads only to sketchy complex concepts.

All the models of concept combination developed by psychologists assume that the processes of concept combination result in comprehensive complex concepts. Thus, Costello and Keane insist on what they call “the semantic richness of conceptual combination”:

People’s interpretations for novel noun-noun phrases are often semantically rich, containing detailed knowledge drawn from various, apparently semantically distant, parts of world knowledge. (Costello & Keane, 2000, 303)

By contrast, the Fast and Frugal Heuristics research program suggests that comprehensive complex concepts might not be needed for many of the purposes that give rise to concept combination. One way to build this idea into a model of concept combination is to assume that a stopping rule governs the production of complex concepts. The production of a complex concept might be interrupted as soon as we have gathered enough information about its extension to fulfill the purpose for which this complex concept was needed.

5. A Toolbox of Simple Heuristics for Concept Combination

Section 4 was mostly critical: we argued that past research on concept combination overlooked the main ideas of the Fast and Frugal Heuristics research program. In the final Section of this chapter, we develop some specific hypotheses about the simple heuristics that might underlie concept combination.

5.1 The Exemplar Heuristic

Many researchers have argued that people's acquaintance with specific members of the extensions of complex concepts plays an important role in concept combination (e.g., Hampton, 1987; Medin & Shoben, 1988; Murphy, 1988; Gray & Smith, 1995; Prinz, 2002). For instance, our acquaintance with some pet fish (e.g., goldfish) explains why we judge that pet fish have few of the properties that we take to be typical of pets and of fish. We concur that people's acquaintance with specific members of the extension of complex concepts plays such an important role, but we contend that the Fast and Frugal Heuristic research program puts this role in a new light.

In Hampton's Composite Prototype Model and in Costello and Keane's C^3 Model, the retrieval of knowledge about specific members of the extension of a complex concept is only one step in the process of concept combination. This knowledge is integrated with other sources of knowledge. By contrast, we propose that whenever retrieving the memory of an instance of a complex concept fulfills the purposes of concept combination, it exhausts the process of concept combination; this memory is simply used as the complex concept. We call this hypothesized process "the Exemplar Heuristic" (for

a similar idea, see Prinz, 2002).⁸ To illustrate, suppose that we are acquainted with a Harvard graduate who is a carpenter and suppose that we are asked whether Harvard graduates who are carpenters are idealistic. According to the Exemplar Heuristic, we retrieve the memory of this individual from long-term memory and we answer positively if our acquaintance is idealistic and negatively if she is not.⁹ The Exemplar Heuristic is a simple process that does not require integrating our knowledge about a specific member of the extension of a complex concept with any other potential source of information about this extension, such as relevant prototypes (e.g., the prototype of a Harvard graduate) or the outcome of causal reasoning. It is also ecologically rational in various environments. For instance, when we are asked generic questions about the instances of a complex concept (e.g., “Are Harvard graduates who are carpenters idealistic?”), answering on the basis of a memory of a specific instance is a reliable strategy, provided that this instance is a typical member (which will usually be true).

As in the case of PET FISH, there is evidence for the use of representations of instances of complex concepts in concept combination. To support our hypothesis about the use of the Exemplar Heuristic in concept combination, however, it is necessary to show that knowledge about specific instances is not combined with other sources of knowledge. One of us (L. Lederer) has recently gathered some evidence consistent with the use of the Exemplar Heuristic in tasks involving the production of complex concepts. In the simple-concept condition, participants were given a list of predicative questions involving a noun and a predicate (e.g., “There is a bird. What is the likelihood that it can fly?”) and were asked to answer using a 10-point scale. In the exemplar condition, participants were given a list of predicative questions involving a noun phrase and a predicate (e.g., “There is an Antarctic bird. What is the likelihood that it can fly?”). The noun-phrases were created by modifying the nouns used in the simple-concept condition and the predicates were identical to the predicates used in the simple-concept condition. In the complex-concept condition, participants were given a list of predicative questions

⁸ In the psychology of concepts, “exemplar” refers to the representations of individual category members (e.g., Medin & Schaffer, 1978).

⁹ We might also retrieve multiple instances of the complex concept.

involving the noun phrases used in the exemplar condition, but modified by an additional adjective (e.g., “There is a mountain-dwelling Antarctic bird. What is the likelihood that it can fly?”).

Importantly, in the exemplar condition, but not in the complex-concept condition, Lederer expected participants to be able to retrieve representations of instances of the complex concepts expressed by the noun phrases. Furthermore, she expected these representations to constitute counterexamples to the predicative sentences participants had to consider. Thus, it was expected that upon reading the noun phrase “Antarctic bird” in the exemplar condition, but not upon reading the phrase “mountain-dwelling Antarctic bird” in the complex concept condition, participants would immediately think about specific instances of penguins, which would constitute a counterexample to the predicative sentence “Antarctic birds can fly” and lead them to give it a lower likelihood rating. Lower ratings given to sentences in the exemplar condition by comparison to the ratings in the two other conditions would suggest that participants indeed retrieved instances of the noun phrases (Figure 4).

Put Figure 4 about here

By contrast, if participants did *not* retrieve any exemplar in the exemplar condition, but instead decided according to the Selective Modification model or Hampton’s Prototype Combination model, their average rating in the exemplar condition should have fallen between their ratings in the two other conditions. According to the Selective Modification Model, both the modifier ANTARCTIC and the modifier MOUNTAIN-DWELLING should have decreased the weight of the BIRD prototype’s *locomotion* attribute (i.e., the attribute containing *flying*) only slightly, since neither modifier affected it explicitly. This would have led participants to give the exemplar condition item a lower rating than the simple condition item, but a higher one than the complex concept condition item. Similarly, according to Hampton’s Prototype Composition Model, because few Antarctic things or mountain-dwelling things can fly, the weight of the item’s *can fly* property would have

decreased over the conditions. Participants would again have given a lower rating to the exemplar condition item and the lowest to the complex concept condition item.¹⁰

As expected, however, participants judged items in the exemplar condition to be substantially less likely than items in the simple concept condition ($F(1,59) = 82.54, p < .000$, Cohen's $d = 2.73$) but also to be substantially less likely than items in the complex concept condition ($F(1,58) = 35.05, p = .006$, Cohen's $d = 1.39$) (Figure 5). This is clear evidence that people access some exemplar information during the process of concept combination.

Put Figure 5 about here

But does this finding show that people use the Exemplar Heuristic? To show that people use this heuristic, we need to show that they not only access some exemplar information during concept combination, but that they do *not* integrate this information with other bodies of information, such as prototypes. For instance, we need to show that when they judge how likely it is that Eskimo shoes have rubber soles, they do access exemplars of Eskimo shoes, but do not integrate this exemplar information with their prototype of shoes. We propose that the large size of the difference between the exemplar condition and the two other conditions tentatively supports the view that exemplar information is not integrated with other sources of knowledge. If exemplar information and prototype information were integrated, we would have found a substantially smaller effect sizes than the one we found. Thus, we take Lederer's result to be tentative evidence for the reality of the exemplar heuristic.

5.2 *The No-Combination Heuristic*

Suppose that you have to decide whether any of several objects or which of several objects belongs to a class you are not familiar with. For instance, you might have to select a Chinese cabbage from the vegetable section in a grocery store. You do not

¹⁰ Here we neglect the fact that Hampton acknowledges the role of exemplars in concept combination. Since in Hampton's model, exemplars have been only used to account for emergent properties, it is unclear what prediction his model would make in the experiment described here.

already have a simple concept CHINESE CABBAGE. According to the view on concept combination developed in this chapter, to perform the task of concept combination, you would rely on a simple strategy likely to give you a successful result, given the context you are in. That is, your strategy would result in your selecting a Chinese cabbage. We propose that in this kind of context, a simple and efficient strategy would be to retrieve only one of the concepts from long-term memory (in this case, the head concept CABBAGE) and to use this concept to categorize the candidate vegetables. The simple concept then stands as a proxy for the complex one. If we had to decide which of several vegetables was a Chinese cabbage, we might choose the vegetable that was most similar to our concept (maybe to our prototype) of cabbage. If we had to decide whether any of several vegetables was a Chinese cabbage, we might choose the vegetable that was most similar to our concept (maybe to our prototype) of cabbage, provided this similarity is above some threshold. We call the practice of using a concept A as proxy for a complex concept $A \wedge B$ “the No-Combination Heuristic.” The No-Combination Heuristic is simpler than any of the processes of concept combination suggested in the literature because it consists simply in retrieving a concept from long-term memory. It is ecologically rational whenever we have to decide whether an object is an instance of $A \wedge B$ and instances of $A \wedge B$ are similar to instances of A. So far, there is no evidence that people use the No-Combination Heuristic, but we feel that asking participants to speak aloud in tasks where a simple concept A would be a good proxy for $A \wedge B$ might provide evidence for the use of this hypothesized heuristic.

5.3 The Elimination-by-Typicality Heuristic

Suppose that you have to decide which of two objects (x and y) belongs to a class you are not familiar with (the extension of $A \wedge B$). Suppose also that x and y are similar both to A and to B. Thus, you can use neither the Exemplar Heuristic nor the No-Combination Heuristic. An alternative strategy goes as follows. Consider the typical properties of instances of A and of instances of B. Take the property that is the most typical (either of the instances of A or of the instances of B). If one of the two objects does not possess this property, then eliminate it and conclude that the other object is an instance of $A \wedge B$. If both

objects have this property or if none has it, then repeat this process with the second most typical property. We call this strategy “the Elimination-by-Typicality Heuristic.”

For instance, suppose you have to decide whether chess or solitaire is a game that is a sport. If you believe that *involving some competition* is the most typical property of sports and that it is most typical of sports than any property is of games, then you might consider whether chess or solitaire involves a competition. Because the former does, while the latter does not, you judge that chess is a game that is a sport. If you had had to discriminate between chess and tic-tac-toe, then you would have to consider the property of sports or of games with the second highest degree of typicality, because both chess and tic-tac-toe are competitive games.

The Elimination-by-Typicality Heuristic is simple in the sense that in contrast with the processes considered in the psychological literature on concept combination, the process of concept combination is not assumed to result in a comprehensive complex concept. Rather, an instance of the complex concept is merely represented as possessing the most typical property (or the second most typical property, etc.) of the instances of one of the two combined concepts. It is also ecologically rational whenever an instance of an $A \wedge B$ is likely to have the typical properties of the instances of A and of B.

5.4 The Modality-Specific Heuristic

The goal of concept combination is often simply to comprehend a speaker or writer’s meaning in discourse or in a written narrative, as in the “spit-glue” example above. When this is the case, the complex phrase to be comprehended is almost always embedded in a discourse or narrative context that specifies its intended modality—that is, the discourse or narrative context specifies the modality, or perceptual sense, to which its meaning is relevant. In these cases, the cognizer might only represent those properties of the head concept that are related to that modality. For instance, when producing a complex concept HALF GRAPEFRUIT in order to draw a still life, the cognizer might represent the property *pink*, but not the property *sour*. When producing a complex concept HALF GRAPEFRUIT in order to follow a recipe, on the other hand, the cognizer might represent the property *sour*, but not the property *pink*. We call this strategy for producing complex

concepts “the Modality-Specific Heuristic.” It produces sketchy concepts, rather than comprehensive concepts, and it is tailored to specific purposes.

Conclusion

In this article, we have reviewed three of the most important models of concept combination developed in psychology since the 1980s—Smith and colleagues’ Selective Modification Model, Hampton’s Composite Prototype Model, and Costello and Keane’s C³ Model. We have contrasted these models with the insights about cognitive processing that emerge from Gigerenzer and colleagues’ Fast and Frugal Heuristics research program. We have speculated that the processes underlying concept combination might form a toolbox of simple heuristics that are tailored to the specific circumstances in which concepts are combined. We have described several possible heuristics—the Exemplar Heuristic, the No-Combination Heuristic, the Elimination-by-Typicality Heuristic, and the Modality-Specific Heuristic—and we have provided some tentative evidence for the existence of the Exemplar Heuristic. We acknowledge the speculative nature of our proposal, but we believe that the perspective on concept combination presented in this article is worth exploring further.

References

- Connolly, A. C., Fodor, J. A., Gleitman, L. R., & Gleitman, H. (2007). Why stereotypes don't even make good defaults. *Cognition*, *103*, 1-22.
- Costello, F. J., & Keane, M. T. (2000). Efficient creativity: Constraint-guided conceptual combination. *Cognitive Science*, *24*(3), 299-349.
- Costello, F. J., & Keane, M. T. (2005). Compositionality and the pragmatics of conceptual combination. In E. Machery, M. Werning, & G. Schultz (Eds.), *The Compositionality of Meaning and Content, II: Applications to Linguistics, Psychology and Neuroscience* (pp. 203-216). Frankfurt: Ontos.
- Costello, F. J., & Keane, M. T. (2001). Testing two theories of conceptual combination: Alignment versus diagnosticity in the comprehension and production of combined concepts. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*, 255-271.
- Downing, P. (1977). On the creation and use of English compound nouns. *Language*, *53*, 810-842.
- Ellroy, J. (1987). *The Black Dahlia*.
- Gagné, C. L., & Shoben, E. J. (1997). Influence of thematic relations on the comprehension of modified-noun combinations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*(1), 71-87.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *103*, 650-669.
- Gigerenzer, G., Todd, P. M., & the ABC Research Group (1999). *Simple Heuristics that Make Us Smart*. New York: Oxford University Press.
- Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, *109*, 75-90.
- Gray, K. C., & Smith, E. E. (1995). The role of instance retrieval in understanding complex concepts. *Memory and Cognition*, *23*, 665-674.
- Hampton, J. A. (1982). A demonstration of intransitivity in natural categories. *Cognition*, *12*, 151-164.
- Hampton, J. A. (1987). Inheritance of attributes in natural concept conjunctions. *Memory & Cognition*, *15*, 55-71.
- Hampton, J. A. (1988). Overextension of conjunctive concepts: evidence for a unitary model of concept typicality and class inclusion. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*, 12-32.
- Hampton, J. A. (1996a). Conjunctions of visually based categories: Overextension and compensation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 378-396.
- Hampton, J. A. (1997a). Conceptual combination. In K. Lamberts & D. Shanks (Eds.), *Knowledge, Concepts, and Categories* (pp. 133-160). Cambridge, MA: MIT Press.
- Hampton, J. A. (1997b). Conceptual combination: Conjunction and negation of natural concepts. *Memory & Cognition*, *25*, 888-909.
- Hastie, R., Schroeder, C., & Weber, R. (1990). Creating complex social conjunction categories from simple categories. *Bulletin of the Psychonomic Society*, *28*, 242-247.
- Jönsson, M. C., & Hampton, J. A. (2008). On prototypes as defaults. *Cognition*, *106*,

913-923.

- Kunda, Z., Miller, D. T., & Claire, T. (1990). Combining social concepts: The role of causal reasoning. *Cognitive Science*, *14*, 551-577.
- Levi, J. N. (1978). *The Syntax and Semantics of Complex Nominals*. New York: Academic Press.
- Machery, E. (2005). Concepts are not a natural kind. *Philosophy of Science*, *72*, 444-467.
- Machery, E. (In press). *Doing Without Concepts*. New York: Oxford University Press.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, *85*, 207-238.
- Medin, D. L., & Shoben, E. J. (1988). Context and structure in conceptual combination. *Cognitive Psychology*, *20*, 158-190.
- Murphy, G. L. (1988). Comprehending complex concepts. *Cognitive Science*, *12*, 529-562.
- Murphy, G. L. (1990). Noun phrase interpretation and conceptual combination. *Journal of Memory and Language*, *29*, 259-288.
- Murphy, G. L. (2002). *The Big Book of Concepts*. Cambridge, MA: MIT Press.
- Prinz, J. J. (2002). *Furnishing the Mind: Concepts and their Perceptual Basis*. Cambridge, MA: MIT Press.
- Smith, E. E., Osherson, D., Rips, L., & Keane, M. T. (1988). Combining prototypes: A selective modification model. *Cognitive Science*, *12*, 485-527.
- Storms, G., de Boeck, P., Hampton, J. A., & van Mechelen, I. (1999). Predicting conjunction typicalities by component typicalities. *Psychonomic Bulletin & Review*, *6*, 677-684.
- Wisniewski, E. J. (1996). Construal and similarity in conceptual combination. *Journal of Memory and Language*, *35*, 434-453.
- Wisniewski, E. J. (1997). When concepts combine. *Psychonomic Bulletin & Review*, *4*, 167-183.

Figure 1

APPLE			
Attributes		Values	
Color	1	Red	27
		Green	3
		Brown	-
Shape	0.5	Round	25
		Cylindrical	5
		Square	-
Texture	0.25	Smooth	24
		Rough	4
		Bumpy	2

Figure 2

GREEN APPLE			
Attributes		Values	
Color	2	Green	30
Shape	0.5	Round	25
		Cylindrical	5
		Square	-
Texture	0.25	Smooth	24
		Rough	4
		Bumpy	2

Figure 3

VEHICLE
1 Carries people or things
2 Can move
3 Moves along
4 Has wheels
5 Is powered, has an engine, uses fuel
6 Is self-propelled, has some means of propulsion
7 Is used for transport
8 Is steered, has a driver controlling direction
9 Has a space for passengers or goods
10 Moves faster than a person on his own
11 Man-made

Figure 4

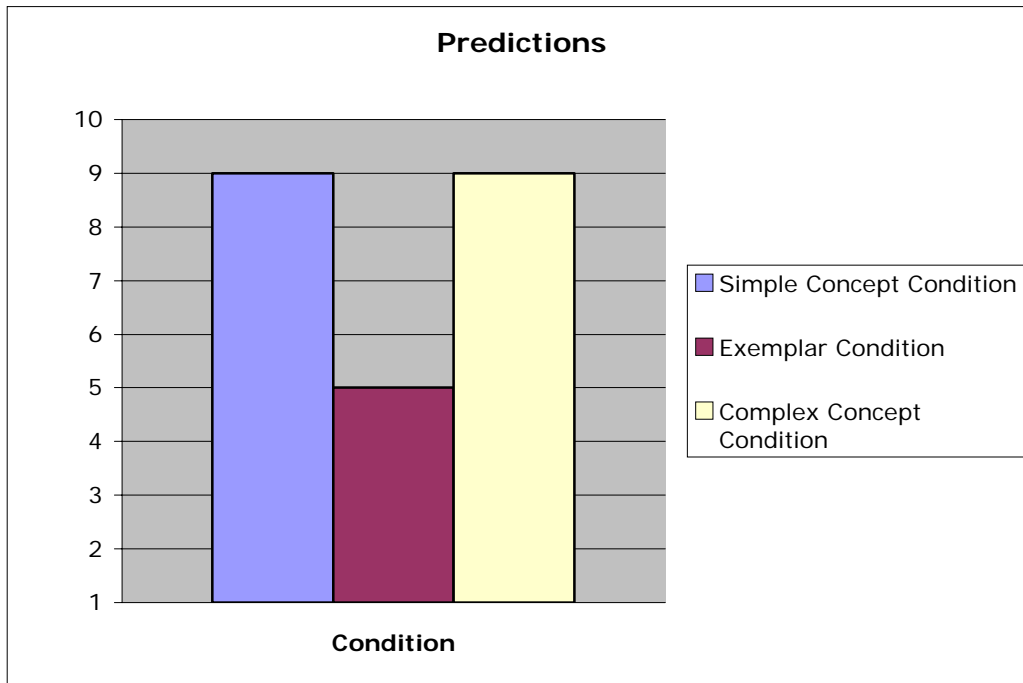


Figure 5

