

Semantics, Concepts, and Meta-cognition: Attributing Properties and Meanings to Complex Concepts

Christina L. Gagné and Thomas L. Spalding

Abstract Lieber (Morphology and lexical semantics. Cambridge University Press, Cambridge, 2004; A lexical semantic approach to compounding. In: Lieber R, Štekauer P (eds) The Oxford handbook of compounding, Oxford University Press, Oxford, 2009) suggests that semantic-encyclopedic knowledge plays an important role in determining the meaning of a compound. In this chapter, we discuss psycholinguistic research demonstrating that conceptual knowledge influences the interpretation of compound words and noun phrases both in terms of a relation-based gist interpretation (e.g., *a green apple is an apple that is green*) and in terms of specific content (e.g., *green apples are sour*). We present evidence indicating that judgments about the properties of modified nouns are affected not only by the content of the constituent concepts but also by meta-knowledge about how sub-categories relate to categories.

Keywords Psycholinguistics • Compounds • Semantics • Properties • Concepts

1 The Challenge of Compositionality

Linguistic morphology focuses on the structure of complex words and word forms. Thus, the issue of meaning is connected to the question of how the meaning of the complex word is related to the constituents. Morpheme-based approaches to this solving this question (e.g., Lieber 2004) use the meaning of the components to specify the meaning of the whole. This is the perspective that we adopt in the current paper – namely, we examine how compound word meaning can be constructed from parts. Our aim is to examine the role of conceptual knowledge in property inclusion in English compound words and to evaluate the extent to which the use of such knowledge is content specific versus abstract meta-knowledge. Property inclusion refers to how properties become associated with concepts. For example, it refers to how people decide whether the property *sour* applies to the

C.L. Gagné (✉) • T.L. Spalding
Department of Psychology, University of Alberta, Edmonton, AB, Canada
e-mail: cgagne@ualberta.ca; spalding@ualberta.ca

concept GREEN APPLE. The question of property inclusion is central to the issue of compositionality and theories of conceptual combination (the process whereby new concepts are formed by combining concepts) differ in terms of how inclusion occurs.

In this chapter, we begin by discussing the challenges of compositionality for linguistic and cognitive theories and provide an overview of three approaches to explaining property inclusion. Next, we present empirical evidence that examines factors that influence property inclusion.

Semantic compositionality poses a challenge in that the meaning of a complex word, although based on the constituents, often contains information that is not directly inherited from the meanings of its constituents. Discussions of word formation often make the valid point that the meaning of derived words cannot be fully compositional due to lexicalization (Dressler 2005; Bauer 1983). Indeed, recently Libben (2014) has emphasized that: “It seems, however, extremely rare that we can say that the meaning of a compound is determined by the meanings of its constituents.” Nonetheless, some aspects of a complex word’s meaning seem to be compositional, though other aspects are not. Lieber (2004, 2009) addresses this issue by suggesting a division between the compositional part of the semantic representation of complex words, which she calls the “skeleton,” and the non-decomposable “semantic-encyclopedic” part, which she calls the “body”.

The semantic skeleton is hierarchically arranged and is based on semantic features that distinguish ontological categories, and time, space and quantity. For example, the skeleton for the word *cookbook* includes the features [+material], [+inanimate] and an R argument which indicates that it is an artifact. The body includes perceptual and cultural knowledge. For example, the body of the word *cookbook* includes knowledge that cookbooks contain recipes. According to this view, the meanings of compound words are determined by the interaction of the skeleton of the bases (i.e., the simple lexemes that make up the complex lexeme). To illustrate, the compound *dogbed* consists of the co-indexing of the R arguments (see Williams 1981 for a discussion of referential arguments) of the skeletons of the simple lexemes *dog* and *bed*. The semantic characteristics of the head override the characteristics of the nonhead such that *dogbed* refers to a type of bed not to a type of dog.

Compositionality is also relevant for psychological theories of concepts and some theorists (Fodor and Lepore 1996; 1998) have gone so far as to conclude that if a theory of concepts cannot explain how concepts combine, then it does not explain what concepts are. In the psychological literature, information that is attributed to the semantic/pragmatic body is viewed as part of the conceptual system rather than as part of the language system because this information is not specific to communication but also is involved in other cognitive functions such as perception, categorization, and reasoning (see Komatsu 1992; Solomon et al. 1999).

Theories concerning conceptual compositionality differ in terms of the extent to which properties of a combined concept are inherited from the constituents. One end of the spectrum is a direct inheritance view which posits that a combined concept initially inherits all properties from both constituent prototypes and that,

in subsequent processing stages, conflicting properties are removed or altered (Hampton 1988, 1991; Jönsson and Hampton 2008; Hampton et al. 2009). For example, Hampton (1991) posits that a combined concept is formed by creating a union of the sets of attributes of both constituent concepts (see also Smith and Osherson 1984). Thus, the attributes of the concepts PET and FISH are combined to create PET FISH. Thus, PET FISH initially has all the attributes of both PET and FISH. This set of attributes is altered according to a set of constraints and principles. An attribute that is a necessary attribute for at least one parent will be retained and is maximally important. Consequently, *having gills* is maximally important for PET FISH because it is a necessary attribute of FISH (Hampton et al. 2009). Incompatible attributes will be removed or altered (e.g., *breathe air*, a typical attribute of PET is inconsistent with *living underwater*, a typical attribute of FISH). By this view, the properties of the constituents are initially available and then removed if they are incompatible with other properties. Additional properties, called emergent properties, may be added based on direct experience with the combined concepts (e.g., from knowledge of the combined concept) or based on the need to resolve conflicts between the inherited properties. For example, *lives in a cage* is an emergent feature of the concept BIRDS THAT ARE ALSO PETS. In sum, the prototype view assumes that the concepts themselves contain the knowledge that is relevant for deriving combined concepts and that prototypical properties are directly inherited during the initial stage of conceptual combination.

At the other end of the spectrum is a non-inheritance view. Connolly and colleagues (2007; see also Fodor 1994; Fodor and Lepore 1996) propose that property inclusion is the result of pragmatic and knowledge-dependent inferences, rather than the result of properties being directly inherited during conceptual combination. By this approach, conceptual combination is a compositional process operating at the level of the entire concept rather than making use of the internal contents (i.e., the properties) of the constituents. According to this view, the initial interpretation of a combined concept is a compositional reading (e.g., a PET FISH is a FISH and is a PET) and particular properties are not inherited from the constituents, but are acquired through experience with the world. As Fodor and Lepore (1996: 265) illustrate: “the language (/concept) assures you that the prototypical pet fish is a pet and a fish just as the language assures you that the prototypically big ant is big for an ant. After that, you’re on your own”.

By this view, rather than being inherited from the constituent concepts, properties arise from non-compositional processing. To illustrate, the initial combinatorial stage indicates that a green apple is both green and an apple. A subsequent processing stage is used to determine what green apples are like and is not based on green or apple but rather comes from whatever someone has learned about the concept GREEN APPLE. Although some of the inferences that people make about green apples follow from their being green and from their being apples, other inferences (e.g., “. . . that green apples go well in apple pie; that they are likely to taste sour; that there are kinds of green apples that you’d best not eat uncooked, and so forth” Fodor 1994: 109) follow from what you have learned about green apples, not from what you know about GREEN and about APPLE.

A third theoretical approach posits that a combined concept starts not with a copy of the head noun which is then altered, but rather with the construction of a structural description that specifies the way in which the two constituent concepts should be connected (Gagné and Shoben 1997; Spalding et al. 2010). For example, the phrase *chocolate bunny* most likely refers to a bunny made of chocolate, but the phrase *chocolate factory* refers to a factory that makes chocolate (for various discussions of relational structures see, Gagné and Shoben 1997; Levi 1978; Libben 1993). Numerous studies have demonstrated that ease of processing is affected by relational availability (e.g., Gagné and Shoben 1997; Gagné and Spalding 2004, 2009). The relational structure provides a gist-based representation of the combined concept and is used in subsequent processing to guide elaboration of the combined concept (Gagné and Murphy 1996; Gagné 2000; Gagné and Spalding 2011; Spalding et al. 2010); to illustrate, chocolate factories do not melt, whereas chocolate bunnies do.

Thus, by this view, judgments about property inclusion are the result of pragmatic and knowledge-dependent inferences, rather than the result of examining whether a property is or is not contained within the representation of a prototype. That is, when making judgments about whether a particular property is likely to be true of a given concept, people draw on various types of meta-knowledge when reasoning about what a particular subcategory is like (Gagné and Spalding 2011). Two types of meta-knowledge are particularly relevant. People use knowledge that an XY is usually a type of Y and that an XY and Y share common properties but that XY is also in some way distinct. The application of these types of meta-knowledge explain why, when making judgments about whether an XY has property A, people are willing to attribute property A (which is a property that is generally true of Y's) to the subcategory XY but this willingness is reduced due to the knowledge that subcategories are not identical to the higher-level category. This knowledge is not specific to particular concepts, but rather is meta-knowledge that is applied across concepts in general.

Finally, the need to distinguish between the language system and conceptual system in terms of discussing the issue of compositionality is most strongly seen when we consider examples such as *blackboard*. This compound refers to the dark (usually black) surface that is used for writing on with chalk and is synonymous with the term *chalkboard*. The compound *whiteboard* refers to a similar surface but is used for writing on with felt-tip pens. The knowledge that a *blackboard* is a writing surface, not just any board that happens to be black, comes from experience with the word in the context of the object to which it refers. That is, it comes from knowledge of the concept BLACKBOARD, rather than from the concepts BLACK and BOARD. The constituent *black* and *board* help form a gist interpretation that can help identify things and events that correspond to the combined concept. Within the language system, *blackboard*, *chalkboard*, and *whiteboard* are related and used in similar contexts. In word-formation, the relevant objects and situations (e.g., the blackboard, the chalk, the writing, and the classroom) are usually present or highly available in memory, but in word-interpretation these objects and events are often more remote (see Körtvélyessy, Štekauer, and Július Zimmermann 2015, for further discussion of word-formation).

2 Empirical Examinations of Property Inclusion

We will focus on experiments that examine property inclusion, which is the process whereby people attribute properties (e.g., *red*) to a particular concept (e.g., APPLE). Early research on property inclusion has found that properties that are true of the entire phrase but not of the head noun are activated prior to properties that are true of both the phrase and head noun (Springer and Murphy 1992). For example, the property *is white* was available prior to the property *is round* for the combined concept PEELED APPLE. This finding occurred even when the items were presented in a discourse context (Gagné and Murphy 1996). For example, it took less time to verify the statement *peeled apples are white* than to verify the statement *peeled apples are round*. (Gagné and Spalding 2007) examined the impact of conceptual combination on the head noun concept and found that properties that are incompatible with the entire phrase temporarily become less available. It was more difficult to verify, for example, *peaches are sweet* when the statement was preceded by a phrase for which the property was not true (e.g., *unripe peaches*) than when preceded by a phrase for which the property remained true (e.g., *orchard peaches*). These findings suggest that the properties of the constituent concepts (especially of the head noun) are not available prior to properties that are unique to the combined concept. This finding is inconsistent with the theories that take a direct inheritance approach because, in these theories, properties of the noun are initially part of the combined concept's conceptual representation and properties that are unique to the phrase are added during a second processing stage.

2.1 Are Prototypical Properties of the Head Noun Directly Inherited or Inferred?

The question of whether properties are directly inherited during the initial stage of conceptual combination or whether they are inferred during post-compositional processing is a key distinguisher among the theories of conceptual combination. Connolly et al. (2007) examined whether prototypical properties of the constituents (in particular, the head noun) are inherited by the combined concept by determining whether people's acceptance of the truth of a proposition was affected by the degree to which the subject noun is modified. In their experiments, participants rated the likelihood of a prototypical property (e.g., *has webbed feet*) for either unmodified nouns (e.g., *ducks*) or modified nouns (e.g., *quacking ducks*, *baby ducks*, or *baby Peruvian ducks*). The modifiers were selected such that they did not conflict with the property being judged. The data indicated that modified concepts received lower likelihood ratings than the unmodified concepts and that the ratings were lower for atypical modifiers (e.g., *baby ducks*) than for typical modifiers (e.g., *quacking ducks*). Finally, additional modification (e.g., *baby Peruvian ducks*) further reduced the ratings. This reduction in the truth value of the property across the various

contexts was originally called the modifier effect. However, because the effect is not entirely due to the modifier, per se, but, rather, to the entire process of modification, we prefer to refer to this effect as the modification effect.

Numerous studies (Hampton et al. 2011; Jönsson and Hampton 2008; Gagné and Spalding 2011, 2014a) have replicated the modification effect and these results have prompted debate about the nature of the conceptual combination process and, in particular, about whether properties of combined concepts are inherited from the constituent concepts (Hampton 1997; Smith et al. 1988) or whether they arise from post-compositional processing (Connolly et al. 2007; Fodor 1994). Connolly et al. argued that the modification effect provides strong evidence against theories that take a direct-inheritance approach because the finding that ratings are reduced for modified concepts is inconsistent with the claim that prototypical properties of the head noun concept are inherited and retained unless they are inconsistent with properties of the modifier concept. In contrast, Hampton et al. (2011) argued that properties were directly inherited but that participants provided lower ratings for statements about modified concepts due to general uncertainty about what those concepts are like. Hampton et al. (2009) argued that ratings of the properties were due to participants having specific knowledge about the combined concept.

We conducted three experiments to further test predictions of the various theories by using dependent measures that more directly test the critical differences between the two views (Gagné and Spalding 2011). In all experiments, participants performed two tasks. In the first task, participants judged, as quickly as possible, whether a generic statement (e.g., *baby ducks have webbed feet*) was true. Each statement was presented in two parts (the subject, e.g., either *ducks* or *baby ducks*, then the predicate, e.g., *have webbed feet*); we were interested in the time required to respond to the predicate. In a second task, participants rated the likelihood that a property was true of a concept. The advantage of the dichotomous decision is that it provides a more direct measure of property inheritance than does the rating task because it is a measure of whether the property is likely to be true or not and, as such, it does not require a more fine-grained decision about precisely how true the statement is, which might encourage the reliance on pragmatic knowledge. In both tasks, we were interested in the particular response itself, as well as in the time required to make this judgment. As for previous experiments (e.g., Connolly et al. 2007; Hampton et al. 2011), the modifiers were selected such that they did not conflict with the property being judged.

The time required to determine whether a property is true of a given concept (or combined concept) provides useful information for testing the question of whether properties are directly inherited. Response time tests an important distinction between the direct inheritance and non-inheritance views because a crucial difference is the point at which properties of the head noun are added to the combined concept. In particular, the direct inheritance view claims that the property is initially inherited from the head noun and, thus, by this view, deciding whether a statement such as *baby ducks have webbed feet* is true involves consulting the representation for the concept BABY DUCKS. In contrast, the non-inheritance view posits that the initial compositional process results in a structure that indicates the

compositional nature of the combined concept (e.g., *ducks that are babies*) but does not include properties. To decide whether the statement is true, participants must engage in inferential and elaborative processing. Consequently, it should take longer to make property judgments for modified nouns than for unmodified nouns due to the additional processing required to infer properties about combined concepts.

In Experiment 1 (Gagné and Spalding 2011), we found that likelihood ratings of the property (e.g., *have string*) were higher for unmodified nouns (e.g., *kites*) than for modified nouns (*silk kites*). Also, people were more likely to respond yes to the unmodified statements than to the modified statements. Responses to the unmodified noun statements were faster than to the modified noun statements. This was not due to reading time required due to the extra word because this measure of the reading time was only to the predicate, which was identical across the experimental conditions. These findings suggest that properties become available later in processing and are not initially available.

In Experiments 2 and 3, we tested the extent to which the modification effect is influenced by the content of the modifier. Experiment 2 used pronounceable nonwords as modifiers. These modifiers were “content-free” because they did not refer to any known concept. This manipulation mimics the real-world situation in which a person has no knowledge about a particular word, and allows us to examine whether the reduction in ratings for modified nouns was due to participants using specific knowledge about the concepts (as claimed by Hampton et al. 2009). Experiment 3 compared statements using unmodified concepts (e.g., *kites*), known modifiers (e.g., *silk kites*) and content-free (nonword) modifiers (e.g., *cary kites*).

The direct inheritance view predicts that prototypical properties of the head noun will be retained in the combined concept unless those properties conflict with properties from the modifier concept. Therefore, in the case of a combined concept using an unknown (i.e., nonword) modifier (e.g., *fleg duck*) prototypical properties should be retained. The judgments for the modified and unmodified nouns should be the same according to this view because without specific knowledge about the modifier, there would be nothing to conflict with the prototypical properties of the head noun and also because the greater the uncertainty the more likely the participants should be to default to the stereotype.

In contrast, the non-inheritance view claims that property inclusion requires knowledge of the combined concept and, therefore, predicts that people should be reluctant to attribute properties to a combined concept containing an unknown modifier because they do not have any real-world experience with the combined concept that can be used to determine which properties are appropriate; Fodor and Lepore (1993: 24) argue that knowing whether a brown cow is dangerous is based on knowledge about the concept BROWN COW, not about the concepts BROWN or COW and that inferences are “. . . determined by ones “real world” beliefs about brown cows.” Because participants will have no real-world experience with the combined concepts denoted by phrases using the nonword modifier (e.g., *cary kite*), they would have no basis for determining whether a particular property is likely to be true. Thus, according to the non-inheritance view the ratings for the modified phrases should be much lower than ratings for the unmodified phrases.

The pattern of data from Experiment 2 which used nonword modifiers was identical to the pattern found from the experiment involving known modifiers. Likelihood ratings were higher for the unmodified concept (e.g., *ducks have webbed feet*) than for the modified concept (e.g., *chonk ducks have webbed feet*). Also, participants were more likely to indicate that the statement was true when judging unmodified nouns than modified nouns. Finally, the time to judge the predicate was longer for modified nouns than for unmodified nouns. Experiment 3 (which involved a direct comparison between known modifiers and content-free (nonword) modifiers) replicated this finding; as in previous experiments, people were less willing to attribute properties to modified nouns than to unmodified nouns and were more reluctant to attribute properties to the content-free modifier items than to the known modifier items. In terms of response time, the two modified conditions did not differ.

In sum, the data reported in Gagné and Spalding (2011) are incompatible with predictions from both the direct inheritance and non-inheritance theoretical approaches. Given the lack of knowledge about the modifier, the inheritance view claims that participants should have defaulted to the prototype. That is, properties of the head noun should have been inherited by the combined concept because there were no properties of the modifier that would have conflicted with them. The data is also inconsistent with the non-inheritance view's claim that inferences about properties require real world knowledge about the combined concept; the data showed that people are willing to attribute properties of the head noun to the combined concept. The response time data also indicate that the properties are not directly inherited during the initial stage of conceptual combination; it took more time to make judgments about modified nouns than about unmodified nouns. If the representation of the combined concept was formed by starting with the representation of the head noun, then there should have been no processing time difference for the modified and unmodified nouns because the property would have been part of the combined concept when participants read and understood the phrase.

Instead, the results provide evidence of an inferential process that operates at the level of logical forms or structures, which are content-free, as well as operating on the content of the head noun category. In this framework, properties are inferred after a structural interpretation has been derived. This process draws on meta-knowledge (which is knowledge about modification, rather than about particular concepts). In particular, meta-knowledge about category-subcategory relations creates a bias towards viewing a subcategory as similar to, but distinct from, the category (see Gagné and Spalding 2011). That is, people use two types of meta-knowledge: First, that an XY is usually a type of Y and second that an XY shares common properties with Y but also has distinct properties. To illustrate, when making judgments about BABY DUCK, people assume that it is a subcategory of DUCK which creates a bias towards agreeing that the property is true. However, their willingness to agree with the property is reduced due to the assumption that subcategories are not identical to the higher-level category. These two pieces of information are not based on specific content of the constituent concepts (i.e., BABY and DUCK) but

rather on meta-cognition about how phrasal structures map onto concepts and on the relation between subcategories and categories, in general. In other words, these two assumptions are made regardless of the specific content of the phrase. Support for this claim that the inferential process draws on meta-knowledge rather than on relying exclusively on specific knowledge about the particular concepts comes from the finding that participants' willingness to attribute a property was reduced whenever the concept was a subcategory, regardless of whether the content of the concept was known (as in the case of *BABY DUCK*) or unknown (as in the case of *FLEG DUCK*). For example, even when the content of the modifier concept is unknown, people use meta-knowledge to determine that *FLEG DUCK* is likely to be a subcategory of *DUCK* which makes them willing to agree that properties that are generally true of ducks are also true of fleg ducks but to a lesser extent due to the use of the subcategory distinctiveness bias.

This use of metaknowledge is also consistent with the finding that the modification effect does not interact with the centrality of the property; Hampton et al. (2009) found that property verifications of mutable properties (e.g., *are white*), categorical properties (e.g., *are mammals*), and central properties (e.g., *are warm-blooded*) were lower for modified concepts (e.g., *inedible lambs*) than for unmodified concepts (e.g., *lambs*). The reduction was consistent across all three property types, which is inconsistent with a direct inheritance approach. If properties were directly inherited, then the central and categorical properties should be more immune to the modification effect than mutable ones. Although central and categorical properties received higher ratings than mutable properties (for both the modified and unmodified concepts), the reduction in ratings for the modified versus unmodified concepts did not differ for the different property types.

2.2 *Is the Modification Effect Due to Uncertainty?*

In a second series of experiments (Gagné and Spalding, 2014a), we examined whether the modification effect is due to uncertainty. Hampton et al. (2011) proposed that properties are initially inherited but that people reduce the ratings due to general uncertainty. They use the example *albino crocodile* and note that because this concept is unfamiliar, people treat the concept "... with a cautious (if not healthy) suspicion. Not only do they doubt whether albino crocodiles have tails and four legs, they are also inclined (to an equal extent) to doubt whether they are reptiles at all" (ibid.: 245).

Theoretical approaches in the literature make opposite predictions about the impact of reducing uncertainty. According to Hampton et al. (2011), reducing uncertainty should increase people's willingness to attribute properties of the head noun concept to the modified concept (i.e., the modification effect should be reduced). In contrast, Gagné and Spalding (2011) posit that reducing uncertainty by providing information about the relational structure should increase the modification effect relative to a situation in which the relational information is implicit. Reducing

uncertainty should increase the modification effect because explicitly providing the relational structure emphasizes that the phrase refers to a particular subcategory which increases the participants' bias towards creating distinctive subcategories. This bias leads people to infer that a property that is generally true of the head noun concept is less true of the combined concept. To test these competing hypotheses, we examined the impact of reducing three sources of uncertainty: Uncertainty about the relational structure, uncertainty about how the modifier contributes to the relational structure, and uncertainty about category membership.

Experiments 1 and 2 (Gagné and Spalding 2014a) examined whether increasing certainty about the relation structure affects the size of the modification effect. The relational structure was either implicit (e.g., *silk kites*) or explicit (e.g., *kites that are made of silk*). The data from Experiment 1 showed that participants were less willing to infer that what is true of the head noun category (e.g., *have strings*) is true of the subcategory when the relation was explicit than when the relation was implicit. Experiment 2, in which modifiers were pronounceable nonwords, was conducted to rule out the possibility that participants were drawing on specific experience with the combined concept when making their judgments. The data from this experiment were consistent with the previous experiment; participants were less willing to attribute a property to a modified concept than an unmodified concept. To summarize, subcategories in which the relational structure was made explicit (e.g., *birds that are babies*) lead to a bigger modification effect than those in which the relational structure was implicit (e.g., *baby birds*), whether the modifier was a known or unknown word.

In Experiment 3, the relation was explicitly provided and we compared judgments about property inclusion for statements with a known modifier (e.g., *bottles that are for spice are cylindrical*), statements with an unknown (i.e., nonword) modifier (e.g., *bottles that are for brinn are cylindrical*), and statements with an unmodified concept (e.g., *bottles are cylindrical*). This manipulation allowed us to examine whether the modifier makes an additional contribution to the inferential process, above and beyond the role that it plays in the relation selection process. Participants were equally likely to agree that the property was true of the concept when the modified concept had a known modifier as when it had an unknown modifier. However, the likelihood ratings were higher for statements with known modifiers than with unknown modifiers. A similar effect was obtained by Gagné and Spalding (2011) with modifier-noun phrases.

The results from all three experiments argue against the hypothesis that the reduction in people's willingness to agree that a property is true is due to uncertainty. On the contrary, increased certainty about relational structure further reduced people's willingness to attribute a property of the head noun to the combined concept.

Experiment 4 of Gagné and Spalding (2014a) examined the consequence of reducing uncertainty about category membership. Participants were explicitly told that an unknown subcategory is a member of a known category. For example, they were told "brinns are a type of bottle". Then they judged a statement about the subcategory (e.g., *brinns are cylindrical*) or about the category (e.g., *bottles are*

cylindrical). The results showed that even when category membership is made explicit, participants were less willing to attribute properties to the subcategory to the same degree that they attributed the same property to the category. In addition, this experiment showed that the modification effect occurred even when the subcategory is specified by single-name label; there was no modification involved. This finding suggests that the reduction in property verification ratings and verification agreement is a consequence of subcategorization rather than of modification, *per se*.

The idea that property verification for both combined concepts and single-term subcategories involves inferential processing rather than the direct inheritance of properties from the higher-level category is consistent with previous results demonstrating that property verification decisions about even simple concepts are affected by recent context. For example, properties of the head that are inconsistent with a recently viewed combined concept become more difficult to verify in the context of the head noun concept (Spalding and Gagné 2007). Verification times to the statement *apples are red* were longer when preceded by a statement containing a combined concept for which that property is no longer true (e.g., *peeled apples are round*) than when preceded by a statement containing a combined concept for which the property remains true (e.g., *autumn apples are round*). This finding is consistent with the claim that judgments about property inclusion are based on inferential processes occurring during an elaboration stage (Spalding et al. 2010; Gagné and Spalding 2011, 2014a).

Furthermore, the claim that elaboration occurs after the formation of a relation-based interpretation is further supported by the finding that deriving property-based interpretations is more difficult than deriving relation-based interpretations. For example, in Experiment 3 of Gagné (2000), participants read a combination (e.g., *whale boat*) and pressed the spacebar as soon as they thought of a meaning. They then typed in the meaning. It took longer to think of a meaning for items that were interpreted using a property (e.g., *a big boat*) than for items that were interpreted using a relation (e.g., *a boat for hunting whales*). This finding extended to a verification task; participants took longer to determine that a property definition was acceptable (e.g., *a whale boat is a large boat*) than to determine that a relation definition was acceptable (e.g., *a whale boat is boat for hunting whales*). Gagné (2000) concluded that relation-based interpretations are the product of an initial compositional process, whereas the property-based interpretations are derived from relation-based based structures and are the product of a subsequent inferential process.

2.3 Does Meta-knowledge Affect Property Inclusion Judgments of False Properties?

The research discussed thus far has focused on properties that are likely to be true; for example, it is reasonable to assume that both ducks and baby ducks have webbed feet. However, categories and subcategories can contrast in terms either

of the properties they have or the properties they do not have. Thus, to further test our proposal that one of the major causes of the modification effect is the implied contrast between the subcategory identified by the modified concept and the category identified by the unmodified concept, we recently conducted a set of experiments in which we also examined property verification of false properties (e.g., *purple candles have teeth* vs. *candles have teeth*).

The direct inheritance approach predicts that judgments about modified and unmodified concepts should not differ because the false property is not part of the head noun concept and thus, should not be inherited. Furthermore, because the property is false, there is no reason for adding the property to the modified concept (i.e., nothing about either the head or the modifier should suggest adding *teeth* to the representation of candles). In contrast, our approach predicts that the bias to make subcategories and categories distinct will make participants more willing to accept that a false property might be true of a modified concept.

In the first experiment, we manipulated whether the concept was unmodified (e.g., *bottle*) or modified (e.g. *brinn bottles*). Because we are interested in the process of conceptual combination, we used nonword modifiers so that participants could not draw on actual experience with the combined concept. We also manipulated whether the property was true (e.g., *are cylindrical*), possible (e.g., *is soft*) or false (e.g., *have teeth*). Thus, there were six versions based on each head noun: Unmodified and True (e.g., *celery is green*), Modified and True (e.g., *meath celery is green*), Unmodified and Possible (e.g., *celery is soft*), Modified and Possible (e.g., *meath celery is soft*), Unmodified and False (e.g., *celery can dance*) and Modified and False (e.g., *meath celery can dance*). Participants saw one version of each item, and an equal number of items in each condition.

As in our previous experiments on this issue, we used two tasks. In the property verification task participants are provided with either an unmodified (e.g., *celery*) or modified noun (e.g., *meath celery*). After reading the item, participants were presented with a property (e.g., *is green*) and asked to indicate (“yes” or “no”) whether the property is likely to be true of the concept. In the second task, participants were presented with a generic statement and rated on a scale from 1 to 10 how likely the statement is to be true.

We analyzed the data using linear-mixed effects (LME) regression models in which subject and item were entered as random factors, and Modification (i.e., modified vs. unmodified) and Property type (true, possible, untrue) were entered as fixed effects (Baayen et al. 2008; Pinheiro and Bates 2000). When fitting a model to predict the binomial response (i.e., “yes” or “no”), we used a generalized mixed-effects multiple regression model with a logistic linking function. In terms of the rating data, modification and property interacted and thus we conducted separate analyses for each property type. For the True properties, ratings were higher for the unmodified noun statements than for the modified noun statements, $b = 1.52$, $t = 13.37$, $p < .001$, replicating the modification effect. However, both the Possible and Untrue properties led to a very different pattern of data. Ratings were higher for the modified than the unmodified noun statements for both the Possible, $b = -.49$, $t = -4.10$, $p < .001$, and Untrue properties, $b = -.37$, $t = -3.80$,

$p < .001$, indicating an inverse modification effect for these two property types. To illustrate, people provided higher ratings for the statement “meath celery can dance” than for the statement “celery can dance”.

The binary decision data showed the same pattern. Modification and property interacted. Participants were more willing to respond “yes” to the unmodified noun statements than to the modified noun statements for the True properties, $b = .96$, $z = 4.86$, $p < .001$, replicating the standard modification effect. In contrast, for both the Possible properties, $b = -.36$, $z = -2.47$, $p < .05$, and the Untrue properties, $b = -.97$, $z = -2.80$, $p < .01$, participants were more willing to respond “yes” to the modified noun statements (e.g., *meath celery can dance* and *meath celery is soft*) than to the unmodified noun statements (e.g., *celery can dance* and *celery is soft*).

Consistent with the claim that the modification effect arises largely as a result of an expected contrast between the subcategory identified by the modified noun and the category identified by the unmodified noun (Gagné and Spalding 2011, 2014a), we find both the modification effect for highly likely properties and an inverse modification effect for the highly unlikely properties. The bias towards viewing the subcategory as somehow different than the category led participants to decreasing ratings and “yes” responses to modified noun statements when the property was true, but to increasing ratings and “yes” responses to the modified nouns statements when the property was possible or untrue.

In the second experiment, we further tested our claim that property judgments are affected by the implied contrast between a subcategory and a category. In this experiment, we directly stated a likelihood of the property for the unmodified head category, and then asked participants to estimate the percentage of members of either the unmodified head category or the modified subcategory that have the property. The property was a blank predicate, which is a predicate for which the participants have no pre-existing knowledge (Osherson et al. 1990). For example, participants might be told “almost all bottles are cooled in annealing ovens” and then asked “What percentage of brinn bottles are cooled in annealing ovens?”

We used three levels of likelihood: Almost all members have the property, Some members have the property, Almost no members have the property. If the implicit contrast between the subcategory and the category is responsible for the effects seen in Experiment 1 and in previous modification experiments, then we should see a modification effect at the Almost all level, no modification effect at the Some level, and an inverse modification effect at the Almost no level. The data support these predictions. There was a strong effect of the likelihood manipulation ($p < .001$), no main effect of modification ($p = .37$), and a significant interaction ($p = .01$). We therefore analyzed the effect of modification at each level of likelihood. At the Almost All level of likelihood, the percentage estimations were higher for the unmodified version than for the modified version ($b = 2.2$, $t = 3.03$, $p = .002$), replicating the modification effect for true properties. At the Almost No level of likelihood, the percentage estimations were higher for the modified version than for the unmodified version ($b = -1.3$, $t = -2.63$, $p = .01$), replicating the inverse modification effect seen in the false properties in Experiment 1. At the Some level

of likelihood, modification was not a successful predictor variable ($b = .46$, $t = .53$, $p = .63$). Thus, when the property was presented as neither common nor uncommon for the head noun category, modification did not influence the likelihood of the property.

3 Implications and Conclusions

We have discussed several results showing that participants' willingness to attribute a property was reduced whenever the concept was a subcategory, regardless of whether the content of the concept was known (as in the case of BABY DUCK) or unknown (as in the case of FLEG DUCK or BRINNS). Participants used meta-knowledge to determine that the combined concept is likely to be a subcategory of the head noun concept and thus they were willing to agree that properties that are generally true of the head noun are also true of the combined concept, but to a lesser extent due to the use of the subcategory distinctiveness bias. In addition to accounting for the modification effect, the subcategory distinctiveness bias accounts for why people are more willing to indicate that an untrue feature might be more likely for a combined concept than for the head noun concept, as well as for why property ratings and acceptance are lower for subcategories denoted by single-word labels.

The research on property inheritance outlined above raises questions about the appropriate metaphor for the relations between properties and concepts. In particular, the most common metaphor for this relationship is one of containment; the concept somehow contains or is made up of, the properties (e.g., Laurence and Margolis 1999). On this view, it is natural to think of the property verification task as involving the participant somehow accessing, or looking inside, the concept to see if the property is there. This, in turn, leads one to think of the properties being "put into" the combined concept, so that property verification for a modified concept also involves looking inside the modified concept to see if the property is there. Our results, in contrast, point to a view of concepts in which properties are not parts of concepts, but rather things that can be predicated/inferred about concepts (see Prasada and Dillingham 2009; Spalding and Gagné 2013). This view is consistent with a long tradition of philosophical work on the relation between properties and concepts (see Spalding and Gagné 2013, for a discussion of this tradition).

These contrasting metaphors also relate to a possible contrast in the understanding of combined concepts in linguistic and psycholinguistic analyses. Linguistic analyses of the semantics of complex words (e.g., Lieber 2004, 2009) are based on characterizing the semantics of the words and their parts and the possible relations between the sets of semantic characteristics. Hence a containment metaphor is natural for such work, as the key point is to characterize the sets of semantic characteristics. However, it would be just as possible to think of the point of the analysis to be understanding what people are willing to predicate of the words, and furthermore, the choice of metaphor may have little impact, given the goals of the

analysis. A psycholinguistic approach, on the other hand, requires consideration of the way in which the person transitions from the parts to the complex, and here the choice of metaphor for the relationship of concepts and properties can have serious theoretical consequences.

Conceptual knowledge is known to influence the interpretation of compound words and noun phrases in two ways. Conceptual knowledge is critical in providing gist-based interpretations (see e.g., Gagné and Spalding 2014a, b, for an overview) and in influencing specific content of the combined concept, based on the contents of the head and the modifier (e.g., Connolly et al. 2007, found a larger modification effect for atypical than for typical modifiers). In addition, we have shown that the gist-based representation is followed by a reasoning process that makes use of the relational gist structure and contents of the head and modifier, as well as meta-knowledge, to infer the properties associated with the combined concept.

References

- Baayen, Harald R., Doug J. Davidson, and Douglas M. Bates. 2008. Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language* 59: 390–412. doi:[10.1016/j.jml.2007.12.005](https://doi.org/10.1016/j.jml.2007.12.005).
- Bauer, Laurie. 1983. *English word-formation*. Cambridge: Cambridge University Press.
- Connolly, Andrew C., Jerry A. Fodor, Lila R. Gleitman, and Henry Gleitman. 2007. Why stereotypes don't even make good defaults. *Cognition* 103: 1–22. doi:[10.1016/j.cognition.2006.02.005](https://doi.org/10.1016/j.cognition.2006.02.005).
- Dressler, Wolfgang U. 2005. Word-formation in natural morphology. In *Handbook of word-formation*, ed. Pavol Štekauer and Rochelle Lieber, 267–284. Dordrecht: Kluwer.
- Fodor, Jerry. 1994. Concepts: A potboiler. *Cognition* 50: 95–113. doi:[10.1016/0010-0277\(94\)90023-X](https://doi.org/10.1016/0010-0277(94)90023-X).
- Fodor, Jerry A., and Ernest Lepore. 1993. Why meaning (probably) isn't conceptual role. *Philosophical Issues* 3: 15–35.
- Fodor, Jerry A., and Ernest Lepore. 1996. The red herring and the pet fish: Why concepts still can't be prototypes. *Cognition* 58: 253–270.
- Fodor, Jerry A., and Ernest Lepore. 1998. The emptiness of the lexicon: Reflections on James Pustejovsky's the generative lexicon. *Linguistic Inquiry* 29: 269–288. doi:[10.1162/002438998553743](https://doi.org/10.1162/002438998553743).
- Gagné, Christina. 2000. Relation-based combinations versus property-based combinations: A test of the CARIN theory and the dual-process theory of conceptual combination. *Journal of Memory and Language* 42: 365–389. doi:[10.1006/jmla.1999.2683](https://doi.org/10.1006/jmla.1999.2683).
- Gagné, Christina L., and Gregory L. Murphy. 1996. Influence of discourse context on feature availability in conceptual combination. *Discourse Processes* 22: 79–101. doi:[10.1080/01638539609544967](https://doi.org/10.1080/01638539609544967).
- Gagné, Christina L., and 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. doi:[10.1037//0278-7393.23.1.71](https://doi.org/10.1037//0278-7393.23.1.71).
- Gagné, Christina L., and Thomas L. Spalding. 2004. Effect of relation availability on the interpretation and access of familiar noun-noun compounds. *Brain and Language* 90: 478–486. doi:[10.1016/S0093-934X\(03\)00459-0](https://doi.org/10.1016/S0093-934X(03)00459-0).
- Gagné, Christina L., and Thomas L. Spalding. 2007. The availability of noun properties during the interpretation of novel noun phrase. *The Mental Lexicon* 2: 241–260. doi:[10.1075/ml.2.2.06gag](https://doi.org/10.1075/ml.2.2.06gag).

- Gagné, Christina L., and 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.
- Gagné, Christina L., and Thomas L. Spalding. 2011. Inferential processing and meta-knowledge as the bases for property inclusion in combined concepts. *Journal of Memory and Language* 65: 176–192. doi:[10.1016/j.jml.2011.03.005](https://doi.org/10.1016/j.jml.2011.03.005).
- Gagné, Christina L., and Thomas L. Spalding. 2014a. Subcategorization, not uncertainty, drives the modification effect. *Language, Cognition, and Neuroscience* 29: 1283–1294. doi:[10.1080/23273798.2014.911924](https://doi.org/10.1080/23273798.2014.911924).
- Gagné, Christina L., and Thomas L. Spalding. 2014b. Conceptual composition: The role of relational competition in the comprehension of modifier-noun phrases and noun-noun compounds. In *The psychology of learning and motivation*, ed. Brian H. Ross, 97–130. New York: Elsevier.
- Hampton, James 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, James A. 1991. The combination of prototype concepts. In *The psychology of word meanings*, ed. Paula J. Schwanenflugel, 1–116. Hillsdale: Erlbaum.
- Hampton, James A. 1997. Emergent attributes in combined concepts. In *Creative thought: An investigation of conceptual structures and processes*, ed. Thomas B. Ward, Steven M. Smith, and Jyotsna Vaid, 83–110. Washington, DC: American Psychological Association.
- Hampton, James A., Martin L. Jönsson, and Alessia Passanisi. 2009. The modifier effect: Default inheritance in complex noun phrases. In *Proceedings of the 31st Annual meeting of the Cognitive Science Society*, 303–308. Austin, TX.
- Hampton, James A., Alessia Passanisi, and Martin L. Jonsson. 2011. The modifier effect and property mutability. *Journal of Memory and Language* 64: 233–248. doi:[10.1016/j.jml.2010.12.001](https://doi.org/10.1016/j.jml.2010.12.001).
- Jönsson, Martin L., and James A. Hampton. 2008. On prototypes as defaults (Comment on Connolly, Fodor, Gleitman and Gleitman, 2007). *Cognition* 106: 913–923. doi:[10.1016/j.cognition.2007.02.009](https://doi.org/10.1016/j.cognition.2007.02.009).
- Komatsu, Lloyd K. 1992. Recent views of conceptual structure. *Psychological Bulletin* 112: 500–526.
- Körtvélyessy, Lívia, Pavol Štekauer, and Július Zimmermann. 2015. *Word-formation strategies: Semantic transparency vs. formal economy*. In *Semantics of complex words*, Studies in morphology, vol. 3, ed. Laurie Bauer, Lívia Körtvélyessy, and Pavol Štekauer, XXX–XXX. Dordrecht: Springer.
- Laurence, Stephen, and Eric Margolis. 1999. Concepts and cognitive science. In *Concepts: Core readings*, ed. Eric Margolis and Stephen Laurence, 3–18. Cambridge, MA: MIT Press.
- Levi, Judith N. 1978. *The syntax and semantics of complex nominals*. New York: Academic.
- Libben, Gary. 1993. Are morphological structures computed during word recognition? *Journal of Psycholinguistic Research* 22: 535–544.
- Libben, Gary. 2014. The nature of compounds: A psychocentric perspective. *Cognitive Neuropsychology* 31: 8–25. doi:[10.1080/02643294.2013.874994](https://doi.org/10.1080/02643294.2013.874994).
- Lieber, Rochelle. 2004. *Morphology and lexical semantics*. Cambridge: Cambridge University Press.
- Lieber, Rochelle. 2009. A lexical semantic approach to compounding. In *The Oxford handbook of compounding*, ed. Rochelle Lieber and Pavol Štekauer. Oxford: Oxford University Press.
- Osherson, Daniel N., Edward E. Smith, Ormond Wilkie, Alejandro Lopez, and Eldar Shafir. 1990. Category-based induction. *Psychological Review* 97: 185–200.
- Pinheiro, José, and Douglas Bates. 2000. *Mixed-effects models in S and S-plus*. New York: Springer.
- Prasada, Sandeep, and Elaine M. Dillingham. 2009. Representation of principled connections: A window onto the formal aspect of common sense conception. *Cognitive Science* 33: 401–448. doi:[10.1111/j.1551-6709.2009.01018.x](https://doi.org/10.1111/j.1551-6709.2009.01018.x).

- Smith, Edward E., and Daniel N. Osherson. 1984. Conceptual combination with prototype concepts. *Cognitive Science* 8: 337–361. doi:[10.1207/s15516709cog0804_2](https://doi.org/10.1207/s15516709cog0804_2).
- Smith, Edward E., Daniel N. Osherson, Lance J. Rips, and Margaret Keane. 1988. Combining prototypes: A selective modification model. *Cognitive Science* 12: 485–527.
- Solomon, Karen O., Douglas L. Medin, and Elizabeth Lynch. 1999. Concepts do more than categorize. *Trends in Cognitive Sciences* 3: 99–105.
- Spalding, Thomas L., and Christina L. Gagné. 2007. Semantic property activation during the interpretation of combined concepts. *The Mental Lexicon* 2: 25–47. doi:[10.1075/ml.2.1.03spa](https://doi.org/10.1075/ml.2.1.03spa).
- Spalding, Thomas L., and Christina L. Gagné. 2013. Concepts in Aristotle and Aquinas: Implications for current theoretical approaches. *Journal of Theoretical and Philosophical Psychology* 33: 71–89. doi:[10.1037/a0029990](https://doi.org/10.1037/a0029990).
- Spalding, Thomas L., Christina L. Gagné, Allison C. Mullaly, and Hongbo Ji. 2010. Relation-based interpretation of noun-noun phrases: A new theoretical approach. In *New impulses in word-formation (Linguistische Berichte Sonderheft 17)*, ed. Susan Olsen, 283–315. Hamburg: Buske.
- Springer, Ken, and Gregory L. Murphy. 1992. Feature availability in conceptual combination. *Psychological Science* 3: 111–117.
- Williams, Edwin. 1981. Argument structure and morphology. *Linguistic Review* 1: 81–114.

Semantics of Complex Words

Bauer, L.; Körtvélyessy, L.; Štekauer, P. (Eds.)

2015, X, 319 p. 42 illus., 2 illus. in color., Hardcover

ISBN: 978-3-319-14101-5