

Conceptual Combination, Property Inclusion, and the Aristotelian-Thomistic View of Concepts

Christina L. Gagné, Thomas L. Spalding and Matthew Kostelecky

Abstract Understanding how properties are extended to combined concepts is critical to theories of concepts. In human judgments, properties true of a noun (ducks have webbed feet) become less true when that noun is modified (baby ducks have webbed feet), while properties false of a noun (candles have teeth) become less false when that noun is modified (purple candles have teeth). These modification and inverse modification effects have been shown to be extremely robust. Gagné and Spalding (2011, 2014b; Spalding and Gagné 2015) have argued that these effects are driven by expectation of contrast. The current experiment shows that, as expected, the modification and inverse modification effects are unaffected by the normative force with which a property is predicated of the head noun, supporting the expected contrast explanation. The results are discussed with respect to an Aristotelian-Thomistic approach to concepts (Spalding and Gagné 2013).

1 Overview

The human conceptual and language systems are compositional and highly productive. One important mechanism underlying this compositional and productive character is nominal compounding. Nominal compounding involves the combination of two or more free morphemes to form a noun. This process can lead to the construction of new concepts in the conceptual system and new words in the mental lexicon. Consequently, the process provides insight into both word formation and conceptual combination. Combined concepts are typically expressed as either

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

T.L. Spalding
e-mail: spalding@ualberta.ca

M. Kostelecky
Department of Philosophy, University of Alberta, Edmonton, Canada
e-mail: kostelec@ualberta.ca

modifier-noun phrases (e.g., *mountain magazine*) or compound words (e.g., *snowball*) and these linguistic expressions are among the most basic examples of linguistic productivity in that they consist of the combination of two morphological constituents. In this chapter, we provide a brief overview of some psycholinguistic and cognitive research that identifies factors that influence the processing of compound words and modifier-noun phrases (which are generally either noun-noun or adjective-noun constructions). We then turn to the question of how the meanings of these linguistic expressions are elaborated, and present recent research on property verification. We discuss implications of these results for human concepts research, and end by suggesting that a view of human cognition derived from Aristotle and Thomas Aquinas could be helpful in understanding recent research on concepts.

2 Background

Concepts can often be paraphrased by expressing an underlying implicit relation between the concepts (*flu virus* is a virus that causes flu). Previous research has found that relational information plays a central role in the processing of combined concepts (for an overview, see Gagné and Spalding 2014a). The interpretation of combined concepts involves selecting the appropriate relational structure, and factors associated with the use of the relations affect the processing of the combined concept (e.g., combined concepts that share a recently used relational structure and a constituent are processed more easily than those that use a different relational structure, and these effects are true of both novel and known combinations; see Gagné and Spalding 2014a). However, such relational interpretations provide only the gist (i.e., a paraphrase) of the combined concept's meaning (e.g., that a *dog house* is a house for dogs) and do not exhaust our knowledge of the combined concept. Where does the rest of the meaning come from? For example, imagine that someone hears, for the first time, the phrase *dog house* and that they know the meaning of *dog* and *house*. Now, suppose we ask this learner, do doghouses have doorknobs? There are several interesting points to note. First, the gist interpretation *house for a dog*, does not, in and of itself, give the answer to the question. Second, the gist is not unrelated or irrelevant to answering the question. The person might reason that because dogs do not have hands, a doorknob would not be likely. Clearly, though, this is an inference based on the gist, not something that is inherently true of the combination of the concepts *dog* and *house*. Third, to be more explicit, note that *dog* and *house* could also be combined to form a gist something like *house that has a dog*, and this type of doghouse would be highly likely to have doorknobs. The key point for the current chapter is that answering the question requires an inference about the specific combination (and gist).

Our primary research question concerns how people make judgments about whether or not particular properties are true for a given concept. In terms of empirical evidence, we will focus on property verification in combined concepts, which refers to how people determine whether a particular characteristic (e.g., “is heavy”) is true

of a modified concept (e.g., *clay machine*). In addressing this question, we apply an approach in which a relation-based interpretation of the combined concept (Gagné and Shoben 1997; Spalding et al. 2010), along with specific knowledge about the constituent concepts and meta-knowledge, is used in subsequent reasoning processes to infer the content of the combined concept (Gagné and Spalding 2011, 2014b; Spalding and Gagné 2015 ; see also Connolly et al. 2007).

We should note that this “relations and reasoning” approach contrasts with many theories of conceptual combination, which tend to assume that the interpretation of a combined concept is the construction of a set of properties true of the combined concept. These properties are assumed mostly to be drawn from the constituent concepts, with some drawn from experience with the combined concept itself (see, e.g., the selective modification model of Osherson and Smith 1981; Smith and Osherson 1984; Smith et al. 1988 or schema based models such as Murphy 1988, 1990; Wisniewski 1997 or prototype combination models such as Hampton 1987, 1991). In turn, these more feature-based notions of conceptual combination tend to assume a kind of implicit “container” metaphor of concepts (see, e.g., Laurence and Margolis 1999, for a discussion of the container metaphor). Such theories tend to assume that properties are contained in the concepts (or, alternatively, that the concept just is the set of properties, or that the concept is made of the properties), and that property verification is, therefore, a matter of somehow inspecting the concept to see if the property is there. A schema theory of concepts, for example, presents the concept as the structured set of properties (i.e., the schema itself is the concept, and the schema is made of the properties in a certain structured set of relationships to each other). Property verification then, is a matter of somehow seeing whether the property is a part of the schema. The relation and reasoning approach and the more feature-based approaches to conceptual combination therefore treat property verification quite differently. We should note that in this discussion of the various models, we have grouped schema-based and property-based theories together because both approaches assume that properties are parts of concepts (that is, they both use a container metaphor of concepts). However, these two approaches differ in other aspects; for example, property-based theories were derived from propositional logic and schema-based theories were derived from predicate calculus (see Barsalou and Hale 1993). These differences, although important, are not relevant for the purposes of our current investigation.

Recently, a series of research projects have aimed at understanding property verification in combined concepts in order to evaluate competing theories of conceptual combination (Connolly et al. 2007; Gagné and Spalding 2011, 2014b; Jönsson and Hampton 2008, 2012; Hampton et al. 2011; Spalding and Gagné 2015). The most important result of these studies has been the modification effect: modification leads properties that are true of the head noun of the modifier-noun phrase to be less true of the modified concept (e.g., *baby ducks have webbed feet* is judged less true than *ducks have webbed feet*) while properties clearly false of the head are judged less false of the modified concept (e.g., *purple candles have teeth* is less false than *candles have teeth*), despite the fact that the properties were explicitly chosen to be unrelated to the modifier that was used. Indeed, the same effects arise even when

the modifier is a non-word (e.g., *chonk ducks have webbed feet*), and therefore cannot provide any conceptual content to act as the basis for the property verification judgment (e.g., Gagné and Spalding 2014b). Furthermore, the modification effect seems to be very robust over levels of typicality of properties of the head. That is, while features highly typical of the head are generally judged more true of the modified concept than features that are not highly typical of the head, typicality of the property for the head does not seem to provide much resistance to the modification effect. Even very highly typical properties of the head (e.g., the property *edible* for the concept *lamb*) show the modification effect (e.g., Hampton et al. 2011; Spalding and Gagné 2015). Finally, we should note that the modification effect appears to be quite robust over specific verification tasks, including ratings of likelihood of the truth of a property for a category, true/false decisions about the property's relation to the concept and the response times to make those decisions, and estimates of the percentages of category members for which the property is true.

(Gagné and Spalding 2011, 2014b; see also Spalding and Gagné 2015) argue that the modification effect is largely due to participants' expectation that subcategories will be both similar to and dissimilar from their categories, and that this suggests that property verification is more a result of reasoning than of inspecting the contents of the combined concept. Furthermore, they argue that the modification effect strongly suggests that properties are not transferred from constituents to the combined concept, but that instead the properties are predicated about the combined concept as a result of the reasoning that results from the demand for property verification.

Thus, we suggest that the processing of a combined concept initially involves constructing a structural description that denotes the way in which the two constituents are connected and determines the linguistic (morphological) structure in which the lexical representations of the constituents are mapped onto particular morphosyntactic roles. This assumption follows from linguistic theories that also propose a structure-based approach (e.g., Allen 1980; Levi 1978; Libben 1993). Thus, the phrase *cherry pie* begins as a modifier-head structure in which the *cherry* is mapped onto the modifier role and *pie* is mapped onto the head role. This linguistic information is involved in constructing a relation-based structure for the corresponding combined concept. The phrase *cherry pie* corresponds to *pie made with cherries*. Properties can then be inferred for this structure. This aspect of our proposal is consistent with previous data indicating that relational structures were available prior to properties. For example, Eliaser (1994) found that relation-based interpretations (e.g., whale boat = boat for hunting whales) were faster than property-based interpretations (e.g., whale boat = a big boat). Similarly, Gagné (2000) demonstrated that it took longer to judge that a property definition was acceptable than to judge that a relational definition was acceptable.

When inferring properties, participants use the meta-knowledge that XY's are a type of Y (see Bauer 1983) to determine that properties that are true of Y (i.e., of the category) are likely to be true of the subcategory, unless there is specific information that would suggest otherwise. However, participants also draw on meta-knowledge that subcategories are different in some way from the category.

Indeed, the purpose of using a combined concept rather than just the category is usually to refer to a subcategory that is distinct from other members of the category (e.g., Clark and Berman 1987; Downing 1977). This accounts for why previous research (e.g., Connolly et al. 2007, 2011; Gagné and Spalding 2014b; Hampton et al. 2011) has found that a property (e.g., has webbed feet) that is true of an unmodified concept (e.g., *duck*) also is seen as being true of a modified concept (e.g., *baby duck*) but less so. This explanation readily accounts for why a property (e.g., has teeth) that is false of an unmodified concept (e.g., *candle*) is judged as being more true of a modified concept (e.g., *purple candle*), as reported in a series of experiments reported by (Spalding and Gagné 2015). Both results (for the true and the false properties) can be explained by participants reasoning that the subcategory is distinct in some way from the category. Interestingly, the reasoning process appears to be more heavily dependent on meta-cognition and on pragmatic factors than on the specific content of the concepts involved. Consequently, content-free modifiers (i.e., unknown words) also produced a modification effect even though they had no content that could be used to assess specific similarity or dissimilarity with the category. Finally, Spalding and Gagné (2015) argue that if property verification for combined concepts involves such predication following reasoning and judgment, it is possible that such processes are also involved in all property judgments for single concepts, not just combined concepts. This, in turn, suggests a very different, more Aristotelian-Thomistic, approach to the relation between property and concept, rather than the kind of implicit “container” metaphor that is so common in current theories of concepts. In particular, in the A-T approach, concepts are not “made of” properties and they do not “contain” properties, and so one does not verify properties by “looking inside” the concept. We return to this point in the Discussion.

3 Current Experiment

In the current experiment, we investigate whether the modification effect is sensitive to the way in which a test property is initially predicated of the head. Spalding and Gagné (2015) used a procedure in which the likelihood of an unknown property was manipulated for a particular head concept, and participants were then asked what percentage of the head concept’s category members had the property or they were asked what percentage of the members of the modified concept had the property. For example, the participants might be told, “Almost all birds have sesamoid bones. What percent of birds/baby birds have sesamoid bones?” Using this procedure, Spalding and Gagné found robust modification effects: The percentage estimates for properties presented as true of almost all category members were higher for the unmodified than for the modified concepts. Spalding and Gagné also found a robust inverse modification effect: When the property was presented as true of almost no member of the category, the percentage estimate was lower for the unmodified than the modified concept. Spalding and Gagné’s manipulation was

about the likelihood of the property for the unmodified head. However, it is unclear whether properties predicated in ways that are more normative, rather than just about the likelihood of the property, would show the same modification effect (Prasada and Dillingham 2009, for example, have shown that simple statistical likelihood does not account for all aspects of the relation of properties to concepts).

In particular, we were interested in whether properties predicated with varying degrees of normative force would show variation in the size of the modification effect. Investigating this point will allow us to determine the extent to which the inclusion of properties is affected by general meta-knowledge about the parent category (e.g., about the category *bird* when reasoning about a subcategory of birds). We chose three ways of predicating the properties. In the Bare Generic condition, the test property was predicated generically with no modification: *birds have property X*. In the Normative condition, the test property was predicated in a normative way: *normal birds have property X*. In the Typical condition, the test property was predicated in a less normative, but frequentist way: *typical birds have property X*. Thus, the Bare Generic condition should have the most normative force, and should thus result in the highest percentage estimates. The Normative condition relaxes the implied “every” from the generic condition, but maintains some degree of normative force, implying that a bird lacking property X might be abnormal. Thus, one might expect that properties predicated in this condition would have lower percentage estimates than those predicated in the Bare Generic condition. The Typical condition further relaxes the force of the predication, such that a bird lacking property X might be atypical, but not abnormal. Thus, one might expect that the properties here would have the lowest percentage estimates of all the predication conditions. The same kinds of predictions could be made for false features (birds do not have property X, normal birds do not have property X, typical birds do not have property X), but with the reversed order (i.e., properties presented as generically false would have the lowest percentage estimation, then those presented as normative, then typical). In other words, the nature of the predication could be seen as directly affecting the strength of the property for the head concept, but without depending on the participants’ specific conceptual knowledge.

The critical interest for the current project, however, is whether this manipulation might affect the modification effect itself. One might make predictions for the size of the modification effect that parallel the predictions for the level of verification. That is, one might think that there should be little or no modification effect for the Bare Generic predication, a somewhat larger modification effect for the Normative condition, and the largest modification effect in the Typical condition. Again, similar, but reversed effects might be predicted for the false properties, across the predication conditions. On the other hand, given the previous work showing, for example, that even categorical properties (e.g., *lambs are mammals*; Hampton et al. 2011) show a robust modification effect, and in which the major factor affecting the modification effect was the expectation of contrast between the category and the subcategory rather than the specific content of the categories (e.g., Spalding and Gagné 2015), one might expect robust and equivalent modification effects across all three predication conditions.

3.1 Method

Materials and design. 84 head nouns were drawn from Gagné and Spalding (2011). Unmodified (*bottles*) and modified nouns (*safety bottles*) were constructed. An unknown predicate (i.e., a potential property of the head which the participant was not expected to know) was associated with each noun. Three levels of normative force were chosen to manipulate the type of predication for the unmodified noun: Bare Generic (e.g., *bottles are cooled in annealing ovens*), Normative (e.g., *normal bottles are cooled in annealing ovens*), Typical (e.g., *typical bottles are cooled in annealing ovens*). A second set of each type of predication was created to present the same unknown property as false of the unmodified noun (e.g., *bottles are not cooled in annealing ovens*; *normal bottles are not cooled in annealing ovens*; *typical bottles are not cooled in annealing ovens*).

The test items were questions asking what percentage of the modified or unmodified nouns had the property (e.g., What percentage of bottles/safety bottles are cooled in annealing ovens?). Predication type was crossed with Property (True/False) and Modification (Unmodified/Modified) to create a $3 \times 2 \times 2$ design.

Participants responded using a scale with a slider that moved from 0 to 100%. These percentage estimates were the dependent variable in the experiment. The materials were counterbalanced such that each level of predication was tested equally often with each modified and unmodified noun. Each participant saw each head noun on only one trial during the experiment. Thus, each person saw 84 trials.

Participants. One hundred forty-four first-year psychology students participated for partial course credit. All were native English-speakers.

Procedure. The instructions included information about the unknown predicates: “The goal of this research is to understand the conceptual system. In everyday life, we are frequently called upon to make judgments of the likelihood of something being true on the basis of limited information. Consider the following situation: If you were told that mammals require biotin for hemoglobin synthesis, how likely do you think it would be that dogs require biotin for hemoglobin synthesis? Because some concepts are described by one word (e.g., dogs), whereas other concepts are described by phrases (e.g., hunting dogs), in this study you will be asked to make judgments involving both types of concepts.”

On each trial, the participants were given a statement in which an unknown predicate was presented with an unmodified noun and participants were told to assume that the statement was definitely true. They were then asked what percentage of either the unmodified or modified noun would have the property (e.g., What percentage of bottles/safety bottles are cooled in an annealing oven?).

Table 1 Mean (SE) estimations of the percent of category members having the test property by property type, predication type, and modification

Condition	Bare generic	Normative	Typical
	True: percent estimation (SE)		
Unmodified	96.7 (0.41)	88.2 (0.52)	88.3 (0.47)
Modified	92.6 (0.61)	83.3 (0.82)	82.7 (0.83)
	False: percent estimation (SE)		
Unmodified	2.1 (0.34)	13.1 (0.66)	12.3 (0.61)
Modified	4.7 (0.52)	16.3 (0.81)	16.5 (0.79)

3.2 Results

The descriptive statistics for the percentage estimates are presented in Table 1. The data were analyzed using Linear Mixed Effects models, as these allow participants and items to be simultaneously treated as random effects (thus avoiding the “language as fixed effect fallacy”, Clark 1973), among other advantages (for discussion, see e.g., Pinheiro and Bates 2000; Quené and van den Bergh 2004, 2008). Predication condition (Bare Generic, Normative, Typical), Property (True, False), and Modification (Unmodified, Modified) were entered as fixed factors. Interactions involving Predication were tested using Chi-square tests of model fits, while other effects involving only factor comparisons with two values are reported as z-scores, based on the slopes of the predictors in the models. The analyses were conducted using the *mixed* function in Stata 13.

Most critically, there is no three-way interaction among Predication, Property, and Modification, $Chi^2(2) = 3.11$, $p = 0.21$. Thus, there is no evidence that the modification effect differs across the levels of Predication and Property.

Because the modification effect is (as predicted) in opposite directions for the true and false properties, it is easiest to understand the data by separating the data on this factor. When considering only the true properties, there is no interaction between Predication and Modification, $Chi^2(2) = 1.99$, $p = 0.37$, but there are strong effects of both Modification ($z = -5.27$, $p < 0.001$) and Predication, with the Bare Generic condition differing from both the Normative ($z = -11.0$, $p < 0.001$) and the Typical ($z = -11.1$, $p < 0.001$) conditions. There is no evidence that the modification effect is sensitive to Predication type.

When considering only the false properties, there is again no interaction between Predication and Modification, $Chi^2(2) = 1.75$, $p = 0.42$, but there are strong effects of both Modification ($z = 3.08$, $p < 0.002$) and Predication, with the Bare Generic condition differing from both the Normative ($z = 12.0$, $p < 0.001$) and the Typical ($z = 13.0$, $p < 0.001$) conditions. As was the case for the true properties, there is no evidence that the modification effect is sensitive to Predication type.

Finally, it is perhaps important to note specifically that there are robust modification effects in the Bare Generic condition, for both true ($z = -6.49$, $p < 0.001$) and false ($z = 4.44$, $p < 0.001$) properties. Thus, although the Bare Generic condition has numerically slightly smaller modification effects, there is no statistical

evidence that those effects differ from the modification effects in either the Normative or Typical condition, and there is evidence for a very robust modification effect, even in the Bare Generic condition.

4 Discussion

The results indicate that the type of predication does matter for property verification, with the Bare Generic form of predication having the highest percentage estimates for the true properties and the lowest for the false properties. This suggests that participants were treating these statements as being the most canonical in that the statements were conveying information that most members of that category share. Indeed, the Bare Generic condition percentage estimates in the current experiment are very similar to those obtained by (Spalding and Gagné 2015, Experiment 2) when they told participants that “almost all” members of the head noun category had the property and then asked participants to estimate percentages. The task used in Spalding and Gagné was identical to the one in the current experiment, except that the likelihood of the property being true of the category was manipulated by using Almost all, Some, and Almost No, suggesting that the Bare Generic predication encourages a reading that is similar to “almost all”, at least in the context of the other predication types. However, the Normative and Typical types of predication did not differ from each other, for either true or false properties. This suggests that participants were treating both the “normal” and “typical” statements as linguistic hedges, rather than as applying differing normative force (as was the case for the Bare Generic form). Indeed, it might be the case that these conditions simply reminded the participants that there are many abnormal and atypical members of most categories, thus lowering the overall likelihood of the property. Thus, it appears that information provided by the type of predication (Bare Generic, Normative, and Typical) is used to provide a starting point for what percentage of category members are expected to have a given property.

Although the type of predication affected the percentage estimates strongly, it had no impact on the modification effect itself (as indicated by the lack of interaction). Thus, as with high levels of typicality or centrality (Hampton et al. 2011; Spalding and Gagné 2015), presenting the property as generically true (or false) of the unmodified head does not provide any resistance to the modification effect, relative to presenting it as normally or typically true (or false). For example, in Hampton et al. (2011), the modification effect was the same for statements with mutable properties (e.g., lambs are white versus baby lambs are white) and immutable (or central) properties (e.g., lambs are warm blooded versus baby lambs are warm blooded). Interestingly, in the current data, despite the fact that the normative and typical framings of the predications resulted in much lower (higher for false properties) percentage estimates, this did not seem to make those conditions any more susceptible to the modification effect than the bare generic condition, as one might expect if those conditions simply made the participants think that

the categories were more variable with respect to those properties. This result (as well as the lack of interaction in terms of mutable versus immutable properties observed in previous studies) is consistent with the view that the modification effect is primarily due to participants' expectations of partial contrast between categories and subcategories (Gagné and Spalding 2011, 2014b; Spalding and Gagné 2015). That is, because those true properties are still estimated at about 85% for the unmodified head (and false properties at about 15%), they drive an expectation of contrast when the noun is modified, but no more nor less of an expectation of contrast than the bare generic condition. To further illustrate, it appears that inferential processing is drawing more heavily on meta-knowledge about the function of modification rather than on content-specific knowledge about the particular properties and concepts, and consequently, for true properties, participants reduce the estimate for modified concepts relative to unmodified concepts due to the expectation that subcategories differ in some way from categories. An estimate about any statement that has the form XY has property P is reduced (for true properties) or increased (for false properties) relative to statements of the form Y has property P. As discussed in the previous paragraph, the starting estimate (i.e., the estimate for the unmodified concept) can differ depending on the form of the prediction, but the size of the increase or reduction is independent of this starting estimate.

To give a concrete example, when participants are told that *bottles are cooled in annealing ovens*, they assume that most members of the *bottle* category have that property. They draw on this information and the meta-knowledge that an XY tends to have the same properties as a Y when estimating what percentage of *safety bottles* have that property to initially consider a relatively high percentage. This estimate is then lowered because participants also use the meta-knowledge that subcategory members are distinct, in some way, from the category. Interestingly, even though there is nothing inherent in the content of the modifier to conflict with that property (just as there is nothing in the content of *baby* that would prevent a *baby lamb* from being warm-blooded), the bias towards viewing categories and subcategories as partially contrastive led to the lowering of the estimate of a property that was previously said to be true of the category. The application of this same bias explains why participants *increase* the estimate for a property that is said to be false of the category; the property that is not true of most members of the category is seen as the basis of the difference between the subcategory and category and, thus, subcategory members are thought to be more likely to hold this property.

One might expect, on this view, that there should be an interaction of modification with predication type, such that the modification effect would be larger as the likelihood of the property for the head goes down (Hampton, personal communication). That is, in the limiting case, if 100% of the category has the property, then 100% of the subcategory should have it, and as one moves away from the 100% mark for the category as a whole, then there is more "room" for the subcategory to contrast with the category (while allowing the percentages to still work out from a mathematical perspective). Two considerations argue against this, however. First, it is highly unlikely that the participants are calibrating their estimates this finely;

instead they seem to simply apply a more or less equivalent level of adjustment of likelihood in all cases. Also, of course, it is known that people are not particularly good about making sure that probabilities or percentages “add up” properly (see, e.g., Jönsson and Hampton 2006). Second, even if they were calibrating their estimates this finely, there is another effect that would seem to be a countervailing force: (Spalding and Gagné 2015) found that when the property was not true enough of the head category, then there was no modification effect. In short, the expectation of contrast depends on there being an initial expectation that the property either is or is not true of the unmodified concept. When participants were told that some X’s have property P, there was no modification effect (nor was there an inverse modification effect). Generally, of course, the existence of both the modification effect and its inverse assumes that at some level of likelihood, there is a reversal of contrast expectation, and hence there are probably some levels of likelihood where there is no effect. In other words, as likelihood of a property decreases towards 50% or so, the modification effect goes away, and then re-emerges as an inverse modification effect as the likelihood of the property moves further toward 0%. Thus, even if participants were trying diligently to make their percentages add up properly, we should be unlikely to see the kind of interaction suggested above, as this change to expected contrast works in opposition to the interaction described in the earlier part of this paragraph.

4.1 The Aristotelian-Thomistic Approach

As suggested by Gagné and Spalding (2011,2014b) and Spalding and Gagné (2015), the effect of expected contrast on the modification effect suggests a very different kind of property verification process than that envisaged in other research on this topic (Connolly et al. 2007; Jönsson and Hampton 2008, 2012; Hampton et al. 2011). In particular, the results suggest that conceptual combination does not consist of the merging and evaluation of properties of the constituent concepts as the mechanism for creating new concepts. That is, the creation of the concept *farm dog* does not appear to consist of accessing properties of *farm* and of *dog* and then building the concept *farm dog* from these properties. If this were the case, then it is difficult to explain the pattern of results observed in the current experiment and in previously published work on the same topic (such as Spalding and Gagné 2015, which also reports an inverse modification effect for false properties, such as *(purple) candles have teeth*). In particular, it is more difficult to understand why properties that were explicitly said to be not true of the category would be added to the subcategory. Furthermore, it strongly suggests that the process of conceptual combination is not a process of mechanically transferring properties from constituent concepts to the newly combined concept, but instead reflects reasoning at the point of decision (for other arguments against such a property transference mechanism, see Connolly et al. 2007; Fodor 1994, 1998, 2001; Fodor and Lepore 1993, 1996). Indeed, these results are difficult to explain using feature-based

theories of conceptual combination. This is particularly true for the false properties, which would have no reason to arise in the combined concept's representation under any of the theories of conceptual combination. The modification effect also suggests that property verification is not a matter of looking into the concept to see if the property is there. Now, as Spalding and Gagné point out, if one does not do property verification by inspecting the contents of the combined concept, perhaps one also does not make property decisions about any concept by inspecting the contents of the concept. How, then, does one make such decisions?

We believe that the current results, and the previous results of research on the modification effect, point to the need for an alternative framework, based on the incorporation of reasoning processes and meta-knowledge about the role of modification in human communication, that could be used in forming and testing psychological theories of concepts and conceptual combination, and in particular theories of the relationship between properties and concepts. Surprisingly, there is relatively little modern research specifically into the relationship between properties and concepts (see, e.g., Laurence and Margolis 1999). Historically, however, the relation between properties and concepts has been investigated in great detail in philosophy, particularly in the work of Aristotle and Thomas Aquinas. (Spalding and Gagné 2013) have argued that the Aristotelian-Thomistic (A-T) view of concepts might have much to tell us about concepts (see also Spalding and Gagné 2015). They showed that the A-T view is not the same as the modern "classical" view, which was (correctly) ruled out by concepts research in the 1970s and 1980s, and discussed several aspects of modern concepts research that fit surprisingly well with the A-T understanding of concepts. Therefore, we end this chapter with a description of properties and concepts in the A-T view, and the suggestion that taking this philosophical view of concepts seriously could have real benefits for research in the psychology of concepts, and particularly for our understanding of the relationship between properties and concepts.

To begin, it is important to realize that thinkers in the A-T tradition have thought deeply about properties and concepts and about how one reasons about things in the world, and have made many distinctions at levels of granularity that are far finer than those in modern concepts research. Also, it should be noted that making judgments about properties is not limited, in the A-T view, to philosophy of concepts, but also explicitly involves the philosophical domains of logic and reasoning. The discussion here is necessarily brief and introductory, but there is much serious discussion of A-T thinking on properties (see, e.g., Reynolds 2001; Wipfel 2000, Part 2), concepts (for an overview, see, e.g., Mercier 1950a) and in the related judgments people make (see, e.g., Mercier 1950b). Indeed, some of the terminology in common modern use is taken from this tradition, though having lost much of its technical meaning in the transition. *Property* is perhaps the most obvious case. *Property* in modern usage is a part of the concept, and as such is a representation of any characteristic believed to be associated with members of the category denoted by the concept. *Property* in the A-T tradition is quite different. First, *property* is a shorthand for *proper accident*, and a proper accident always inheres in things in the world, not in concepts. *Accidents* contrast with *substances*. Substances, to a first

approximation, are things that exist in themselves, while accidents are things that exist in or ‘inhere in’ substances. Thus, neither substances nor accidents (including proper accidents) primarily exist in representations because accidents inhere in substances and substances are found in the real world. Furthermore, proper accidents contrast with other kinds of accidents, in that proper accidents “flow from” the essence of the substance in which they inhere, while other kinds of accidents do not flow from the essence of the substance. For example, having two eyes flows from the essence of being human, but having brown hair does not. Note that being necessary, in this sense, does not mean that every existing human must have two eyes, or that a human who loses an eye becomes non-human, but rather that it is in the nature of the human to have two eyes, and that a failure to have two eyes is a deprivation of the normal property flowing from the essence. Such deprivations are well-known and well-explained by the A-T view (see Spalding and Gagné 2013, for more discussion and comparison to the so-called classical view of concepts). Thus, properties, in the A-T view, are not parts of concepts. Instead, they are characteristics of things in the world, though one may have a concept of a given property, of course. Properties can be predicated of primary substances (e.g., *wings* can be predicated of individual birds in the world), or of secondary substances (i.e., *wings* might also be predicated of a universal such as *birds*). Also, such predications can be remembered and become habitual, because we have understood how *wings* and *birds* relate. Hence, one need not consciously think through the relationship of *wings* and *birds* anew every time. Importantly, though, it is the act of predicating *wings* of *birds* that becomes habitual. Thus, it is not the case that *wings* must become part of the concept *bird*, even when the predication happens habitually (i.e., in the A-T view, it is not necessary to turn to a conceptual container approach to account for the predication).

A second important characteristic of A-T thinking about concepts and properties is that, unlike most thinking about concepts since at least Locke, the concept is not the thing that the mind considers when it thinks about something. Rather, the thing is what the mind considers, while the concept is simply what allows the mind to think about the thing. One thinks about the external object via the concept, rather than thinking terminating in or being directed toward the concept itself. Of course, the mind can reflexively think about a concept—I can consider my own concept of *bird* and reflect upon what it is to have a concept at all, but that is not what I do when I think about birds; note that this distinction is rather hard to keep clear under a Lockean approach, where to think about birds just is to activate the concept *bird*, and thinking about the concept *bird* is also to activate the concept *bird*. Relatedly, thinking, in the A-T tradition, is sometimes said to involve representations, but this is not quite correct, in modern terms. Instead, thinking in the A-T tradition involves *re-presentations* of the thing to the mind, of the mind *re-presenting* the thing internally. Thus, *re-presentation* (in the verb sense) is primary in the A-T view, whereas a *representation* (in the noun sense) is primary in modern theories. In the A-T view, representation of a thing is an event, an action undertaken by a cognizer through which a thing is known, not typically a thing known in itself. Although this distinction may seem picayune, it is important to note that the early modern (and

continuing) difficulties with radical skepticism and the “problem of the bridge” derive from the claim that the ideas are what the mind grasps—If the mind grasps ideas, how do we bridge from the ideas to the world? What becomes a fundamental problem for many modern accounts of knowledge is simply not an issue on the A-T view of knowledge, since in this account the mind grasps the thing through the concept (or, stated differently, the thing is re-presented through the concept) rather than grasping the concept itself. For Aristotle and, especially, Thomas Aquinas, the mind mirrors or re-presents the external world in a basically reliable way (Feser 2009, pp. 145–149, provides a clear discussion of this point; for a more detailed treatment, see Perler 2001).

A third important characteristic of the A-T view of concepts and properties is that all thinking involves the use of sensory information. In the A-T view, any judgment about a thing, for example, involves not only the concept of the thing, but the *phantasm*, the sensory information associated with the actual things that the concept is “of”. There is, following a famous dictum of Aristotle, no thinking without a phantasm. This necessary involvement of the phantasm is important because it means that there is a constant, fundamental, and completely thorough-going recourse to the senses for any thought whatsoever. The A-T view has a long history of carefully considering how exactly sensory and intellectual information relate and interact in thinking. This history could be very helpful in understanding some of the recent work on embodied concepts and perceptual symbol systems (e.g., Barsalou 1999). Indeed, notions such as the re-enactment of conceptual content via simulation (e.g., Barsalou 2009, p. 1282) accord extremely well with the A-T notion of re-presentation (verb) as opposed to representation (noun). Similarly, the way that (Wu and Barsalou 2009) describe the process of generating features from simulations is very similar to how the A-T view would describe the process of predicating properties via the phantasms (which are also multi-modal and sensory-based). Furthermore, the A-T approach incorporates the imagination, a power of manipulating phantasms, such that sensory information can be retrieved, manipulated, and re-combined in novel ways, as well as common sense (here a technical A-T term), which combines sensory information from multiple sensory modalities. Finally, conceptual thinking in the A-T view involves the intellect, which is the power called upon in formal reasoning, abstract concept formation, mathematics, and other such tasks, as well as a power that can direct the imagination and access the phantasms. Thus, when a person thinks about concepts and properties, all of these various powers come into play.

A fourth important characteristic of the A-T view is that properties (or features) are not more fundamental than substances. Rather, substances are more fundamental than properties, because we say or ‘predicate’ properties of substances and not the other way round. For example, we predicate ‘rational’ of ‘Socrates’ and not ‘Socrates’ of ‘rational’. We would never claim that ‘Rational is Socrates,’ though the obverse clearly obtains, namely ‘Socrates is rational.’ The property (rationality) is predicated of the substance (Socrates). Moreover, we never predicate one substance of another. For instance, we cannot predicate Socrates of Plato in the following judgment: Plato is Socrates. Instead, we predicate various accidents of

substances and these accidents or features are seen to inhere in the substance. This contrasts markedly from most modern views, in which concepts consist of properties. That is, the concept is made of or just is the set of properties. Thus, the properties are more basic than the concept, which has several interesting consequences. One is that if this general view were true, one would expect that categorization would be a matter of identifying the properties of a thing and then using those properties to determine whether this thing is in the category. Indeed, as a general description, this is basically how categorization is understood in most modern theories. However, this implies that the properties should be identified prior to the category. Thus, before one grasps that the thing they are looking at is a bird, they should already have grasped that they are seeing a beak and feathers and wings, etc. Note that it is not that they are seeing visual information that can be interpreted in these ways, but literally that they are grasping those conceptual features. A second implication of this approach is that there should be a clear and simple linkage from the extraction of simple visual features to simple conceptual features to the concept. In short, we should have no discontinuity between research on object recognition and research on concepts and categorization. Yet, research on object recognition rarely uses conceptual features, while research on concepts and categorization assumes conceptual features and rarely even posits the exact connection to the simple visual (or, for that matter, other sensory) features.

Finally, it is important to understand that the A-T view of concepts and properties is embedded within a very complicated and intricate understanding of not only the mind, but of the physical world. Thus, for example, the A-T way of thinking about essences, properties, and concepts derives not from what is convenient for philosophy of mind, but from prior metaphysical commitments about the physical world and how the mind is seen to function (fairly reliably) as a part of that world. That is to say, the A-T approach to accounting for physical change through terms like actuality/potentiality and form/matter serves as the basis for A-T cognitive theory, and this cognitive theory is, in turn, a further specification of those basic metaphysical commitments. An immediate and obvious benefit of this approach, say advocates of A-T cognitive theory, is that it is completely consistent across differing levels of cognitive and metaphysical explanation, from how intellectual concept formation occurs, to other sorts of cognition, to an explanation of the constitution of physical things (including humans), and even to the processes of change as such (e.g., Feser 2014) because it is all part of a general and consistent view of the world.

What is surprising about this is how well the philosophy of mind inherent in the A-T approach fits modern evidence about the mind. (Spalding and Gagné 2013), for example, discuss how the A-T approach fits well with a wide variety of evidence drawn from modern psychological research on concepts, including evidence for prototypes and exemplars, as well as work on generics and essentialism, the relation between human and non-human animals' abilities to deal with categories, and other recent research topics in concepts. We should also note here that (Prasada and Dillingham's 2009) notion of K-properties (i.e., properties that are indicative of a kind of thing) is explicitly Aristotelian, though they do not explicitly make use of

the rest of the A-T approach to cognition. Pustejovsky (1995) incorporates a version of the four Aristotelian causes in his “qualia” structure for lexical items. Nor should we think of the A-T view as being relevant only to the study of concepts or language. Feser (2014) provides a broader discussion of the A-T approach as an appropriate underpinning for science in general, both in the sense that its view of human cognition provides a proper underpinning for the human activity of doing science and in the sense that its metaphysical commitments provide a proper underpinning for science considered as a set of propositions about the world. Indeed, the A-T view may well provide at least as good a philosophy of mind as we can find in modern and contemporary approaches (see e.g., Madden 2013; see also Spalding et al. 2014). At the very least, the consistency between cognitive theory and metaphysical commitments in the A-T view presents a picture that contrasts sharply with modern psychology and philosophy of mind, where there seem to be extraordinarily deep and intractable problems in consistently linking down to metaphysics (Madden 2013, for example, discusses many problems with the relations between modern philosophy of mind and modern metaphysical approaches; see also Feser 2006, 2008, 2009).

4.2 Applying the A-T View to the Modification Effect

The applicability of the A-T view to cognitive science is seen acutely in what we have been discussing throughout this article: property verification judgments and subsequent modifications made to those judgments. To show this applicability, it is perhaps best to begin with a general explanation of how a property verification judgment is made (in its strictest sense) on the A-T view (i.e., through what Aristotle and Aquinas call *episteme* or *scientia*, respectively), and then present how the same approach might account for the sort of property verification results encountered in our experiment.

On the A-T view, property verification in its strictest sense occurs in virtue of understanding how a given property is explanatorily useful in knowing the nature of the thing in which the property is found. In other words, in the A-T view, to do property verification is to verify that a characteristic is, in fact, a property of the thing (i.e., that it is a characteristic that flows from the essence of the thing). To put the matter concretely, we can ask, ‘How does one decide whether birds have wings?’ but with the understanding that we are actually asking how one decides whether *wings* is a proper accident (flowing from the essence) of *bird*. On the A-T view, property verification in this strict sense will occur when we understand why the property regularly occurs in the kind of thing of which we are predicating the property. Note that in the A-T view, much of what we “know” about the world, even though a reliable guide for action, etc., is not actually known in the strict sense of *episteme* or *scientia*. We return to this point below.

Now, as we have pointed out above, judgments are not made by looking in the concept to see if there is a property there, but rather through a judgment about the

thing to be known (thus, the distinction between a concept as what the mind grasps vs. that by which the mind grasps the thing to be known, is quite critical to a proper understanding of how property judgments are made). In the A-T view, our knowledge begins in a confused and indeterminate manner and grows more specific as we encounter a thing regularly (if we choose to try to understand). As we regularly encounter birds, we may notice wings, beaks, feathers, etc., and notice that other animals for which we have refined concepts do not have these characteristics. We do not yet know what these characteristics are, nor what they are for, but they may make the object of inquiry stand out with respect to other animals. Perhaps after further inquiry we come to see what these features are for and notice that many or even all birds that we have encountered have these features. Once we have this knowledge (i.e., what these features are for and a wide, perhaps, universal distribution), wings, beaks, and feathers will become part of the causal story of birds, and we can understand why birds have wings (even, with much inquiry, flightless birds). At this point, we have recognized and predicated a proper accident, a property that we think falls to all birds because we see that the nature of ‘bird’ is at least partially, but essentially, explained by the property ‘wing’. We thus engage in a process of developing more precise knowledge as we become more familiar with a given object. Ultimately, in the A-T view, this process can (but does not always) result in a definition, in which the thing can be defined, in terms of its nature, via its genus and specific difference (such that the human being can be defined as a type of animal, specifically a rational animal; note that genus and species are used here as technical terms, but not in their modern biological technical senses).

How, then, would the A-T view explain the sort of property verification that we tested in our experiment, where we are clearly not asking for property verification in the strictest A-T sense? That is, how does one, on the A-T view, account for the judgment that ‘birds have sesamoid bones,’ and, more pressingly, then account for a modification effect when test subjects were told to apply ‘sesamoid bones’ to a newly combined concept, like ‘baby bird’? These examples clearly differ from the property verification in the strict sense described above, and for at least two reasons: first, our test subjects do not know the meaning of the predicate (in this case, ‘sesamoid bones’), whereas by the end of investigation the inquirer, by definition, knows what the predicate means in what is described above. Second, A-T property verification in the strict sense is meant to result in a ‘categorical’ judgment, like ‘*All* birds have wings,’ which applies universally to all things in the relevant category (note, however, our comments above that terms such as “universal” or “necessary” do not actually mean that every existing example will have the properties—there is still the issue of deprivation of such properties—rather it means that such properties are seen to flow from the essence). By contrast, in our experiment, our test subjects are operating on distinctions of ‘generic’, ‘normal’, and ‘typical’, none of which claims that the property universally applies. With those differences in mind, we think that the A-T approach to cognitive theory is consistent with our empirical results regarding reasoning about combined concepts.

An essential part of this story is to account for how and why reason functions as it does when applied to combined concepts, which plays to the strengths of the A-T

view. While the A-T view of cognition has a robust (proposed) philosophical description of the mechanics of cognition (see, e.g., Brennan 1941; Kosteletzky 2014), from how sensory data is initially perceived in a phantasm to an intellectual judgment about the thing, it does not rest solely on those mechanics. Instead, there are different sorts of knowledge and different kinds of inquiry that supervene on the process that results in the conceptualization of the thing. The ‘highest’ form of knowledge (‘highest’ because it is most certain) is simply ‘understanding’. In this form of knowledge, one’s assent is forced, as it were, and there is no process of reasoning required for assent. A classic example is ‘A whole cannot be smaller than one of its parts.’ If I have understood the terms at play, I cannot force myself to assent otherwise, since it is impossible to conceive of a whole being smaller than one of its parts, and, importantly, the truth of the statement is immediately apparent; reasoning and inquiry do not occur. Below this sort of knowledge is ‘science’, which is the so-called strict sense of property verification discussed above. ‘Science’ deals in necessary knowledge, with conclusions that cannot be otherwise than they are, but reasoning is clearly employed because ‘science’ follows the pattern of a demonstrative syllogism. For example, if I know that all humans are rational and that all Greeks are human, I know that all Greeks are rational.

Below this sort of inquiry is ‘dialectics’, which also makes claims that are meant to be universal, but the premises are only taken as likely true, because the premises are acquired or accepted from “the received opinion of the wise and the many” (Aristotle, *Topics* I). This basic aspect of dialectics dovetails nicely with how our research participants received their knowledge that ‘bottles are cooled in annealing ovens’ or that ‘birds have sesamoid bones.’ We told them, and they took this received opinion as true. Nevertheless, on the A-T view they neither ‘understand’ the proposition nor have an account of the ‘reason why’ the proposition is true. It is, at best, ‘probable knowledge’ (i.e., what Aristotle claims is the product of dialectics). Yet, nearly all our participants assented that these properties were true of the head, as can be expected given that we (“the wise and the many”, at least in the context of our experiment!) asserted the truth of these propositions.

As we modified the concept, a modification effect comes into play, and fewer subjects assent to the proposition involving the modified concept. This can be accounted for by way of the relationship between reasoning and assenting to a proposition, because assent, for Thomas Aquinas (he discusses issues of assent directly, whereas Aristotle does not treat it as precisely), occurs more firmly and immediately if the premises themselves are more certainly known. If the premises are not firmly known, more reasoning is required to get to the point of assent (see Thomas’ Disputed Questions on Truth, q. 14, art.1). Even then we are going to hold less firmly the proposition to which we have assented.

On this account of how reasoning and assent function together, it is hardly surprising that fewer people will assent to a proposition with a modified concept, especially if they took the truth of the original proposition on trust. There are two reasons for this. First, of course, we clearly have only probable reasoning in moving from the truth of the property for the head to the truth of the property for the combined concept, and this, according to Thomas, should result in lower confidence

in the outcome of that reasoning. This aspect of reasoning to assent would apply to any proposition that involved probable reasoning (i.e., the strength of the conclusion is going to be lower than the strength of the premise whenever the reasoning involved is only probable, unless direct knowledge of the conclusion, or other knowledge, can intervene).

Second, and more specific to the modification effect per se, the structure of the Thomistic understanding of category and subcategory should give rise to the kind of expectation of similarity and contrast suggested by Gagné and Spalding (2011, 2014b) and Spalding and Gagne (2015). That is, as noted above, in the Thomistic view a subcategory is distinguished from a category via a specific difference (i.e., something that differentiates this species from its genus, as in the human being defined as a rational animal). Hence, one should expect that there is some difference (and, likely, an essential difference) when a subcategory is being posited. Of course, a person can posit a purely accidental or even meaningless subcategory, but that should be a relatively rare event and should not much affect one's expectation that a posited subcategory would differ from the category in meaningful way. Even in the case of a purely accidental subcategory (in A-T terms, i.e., a subcategory that differs primarily in non-proper accidents), there will often be reasons for the non-proper accidental difference that will be related to the essence of the category. Consider, for example, the subcategory of birds that live in Edmonton, Alberta, Canada. Although there is no particular essence that relates to this rather artificial and accidental subcategory, there are differences between the birds that can live in this location and those that cannot, and these differences would be very likely related to the essences of those various types of birds. Hence, one might expect that this group may differ from birds in general in essential ways, even though this grouping is accidental. But, in many, if not most, cases of known subcategories there would be an explicitly essential difference—the essence of a subcategory just is not, generally, the same essence as the category. Furthermore, because multiple accidents (both proper and otherwise) will either flow from, or co-occur with, a difference in essence, there is good reason for the participant to expect contrast on various characteristics, but also to expect that overall, there should still be many similarities (because much that is essential, as well as proper and accidental, is in both the category and subcategory). In a nutshell, this is how (Gagné and Spalding 2011, 2014b; Spalding and Gagné 2015) account for the modification effect. It is not inherently about modification or conceptual combination, but instead is about the expected relationship between the category and the subcategory. That expected relationship of similarity and contrast enters into the reasoning process that the person goes through in evaluating the likelihood of the truth of the predication of the property to the subcategory.

5 Conclusion

The current experimental results are consistent with a view of the modification effect in which the primary driver is the expectation of similarity and contrast between category and subcategory, as suggested by (Gagné and Spalding 2011, 2014b; Spalding and Gagné 2015). In particular, and contrary to what would be expected given feature-based theories of conceptual combination, the modification effect seems to be insensitive to the normative force with which a property is predicated of the head noun concept, though the overall level of attribution of the property is quite sensitive to the force of the predication. This view, in turn, is consistent with an approach to concepts deriving from the Aristotelian-Thomistic tradition (Spalding and Gagné 2013). We argue that taking the A-T view seriously has much to offer current research on concepts and might provide a fruitful basis for the development of a psychological theory of conceptual combination and of concept processing, in general.

References

- Allen, M. R. (1980). Semantic and phonological consequences of boundaries: A morphological analysis of compounds. In M. Aronoff & M. L. Kean (Eds.), *Juncture* (pp. 9–27). Saratoga, CA: Anma Libri.
- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22, 577–660.
- Barsalou, L. W. (2009). Simulation, situated conceptualization, and prediction. *Philosophical Transactions of the Royal Society*, 364, 1281–1289.
- Barsalou, L. W., & Hale, C. R. (1993). Components of conceptual representation: From feature lists to recursive frames. In I. van Mechelen, J. Hampton, R. Michalski, & P. Theuns (Eds.), *Categories and concepts: Theoretical views and inductive data analysis* (pp. 97–144). San Diego, CA: Academic Press.
- Bauer, L. (1983). *English word-formation*. Cambridge, UK: Cambridge University Press.
- Brennan, R. E. (1941). *Thomistic psychology: A philosophic analysis of the nature of man*. New York: Macmillan.
- Clark, H. H. (1973). The language-as-fixed-effect fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior*, 12(4), 335–359.
- Clark, E. V., & Berman, R. A. (1987). Types of linguistic knowledge: Interpreting and producing compound nouns. *Journal of Child Language*, 14, 547–567.
- Connolly, A. C., Fodor, J. A., Gleitman, L. R., & Gleitman, H. (2007). Why stereotypes don't even make good defaults. *Cognition*, 103, 1–22.
- Downing, P. (1977). On the creation and use of English compound nouns. *Language*, 53(4), 810–842.
- Eliaser, N. M. (1994). Distinct processes in conceptual combination. *Unpublished master's thesis, Northwestern Univ., Illinois*.
- Feser, E. (2006). *Philosophy of mind: A beginner's guide*. Oxford: Oneworld.
- Feser, E. (2008). *Last superstition*. South Bend, IN: St. Augustine's Press.
- Feser, E. (2009). *Aquinas: A beginner's guide*. Oxford: Oneworld.
- Feser, E. (2014). *Scholastic metaphysics: A contemporary introduction*. (Editiones Scholasticae: Vol. 39). Heusenstamm: Editiones Scholasticae.
- Fodor, J. (1994). Concepts: A potboiler. *Cognition*, 50, 95–113.

- Fodor, J. A. (1998). *Concepts: Where cognitive science went wrong*. New York, NY: Oxford University Press.
- Fodor, J. A. (2001). Language, thought and compositionality. *Mind and Language*, 16, 1–15.
- Fodor, J. A., & Lepore, E. (1993). Why meaning (probably) isn't conceptual role. *Philosophical Issues*, 3, 15–35.
- Fodor, J. A., & Lepore, E. (1996). The red herring and the pet fish: Why concepts still can't be prototypes. *Cognition*, 58, 253–270.
- Gagné, C. L., & Shoben, E. J. (1997). Influence of thematic relations on the comprehension of modifier-noun combinations. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 23, 71–87.
- Gagné, C. L. (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(3), 365–389.
- Gagné, C. L., & Spalding, T. L. (2011). Inferential processing and meta-knowledge as the bases for property attribution in combined concepts. *Journal of Memory and Language*, 65, 176–192.
- Gagné, C. L., & Spalding, T. L. (2014a). Conceptual composition: The role of relational competition in the comprehension of modifier-noun phrases and noun-noun compounds. In B. H. Ross (Ed.), *The psychology of learning and motivation* (Vol. 59, pp. 97–130). New York: Elsevier.
- Gagné, C. L., & Spalding, T. L. (2014b). Subcategorisation, not uncertainty, drives the modification effect. *Language, Cognition and Neuroscience*, 29(10), 1283–1294.
- Hampton, J. A. (1987). Inheritance of attributes in natural concept conjunctions. *Memory & Cognition*, 15, 55–71.
- Hampton, J. A. (1991). The combination of prototype concepts. In P. J. Schwanenflugel (Ed.), *The psychology of word meanings*. Hillsdale, NJ: Erlbaum.
- Hampton, J. A., Passanisi, A., & Jönsson, M. L. (2011). The modifier effect and property mutability. *Journal of Memory and Language*, 64, 233–248.
- Jönsson, M. L., & Hampton, J. A. (2006). The inverse conjunction fallacy. *Journal of Memory and Language*, 55, 317–334.
- Jönsson, M. L., & Hampton, J. A. (2008). On prototypes as defaults (comment on Connolly, Fodor, Gleitman, and Gleitman, 2007). *Cognition*, 106, 913–923.
- Jönsson, M. L., & Hampton, J. A. (2012). The modifier effect in within-category induction: Default inheritance in complex noun phrases. *Language and Cognitive Processes*, 27, 90–116.
- Kostelecky, M. (2014). Not induction's problem: Aquinas on induction, simple apprehension, and their metaphysical suppositions. In P. Biondi & L. Groarke (Eds.), *Shifting the paradigm: Alternate perspectives on induction*. Berlin: De Gruyter.
- Laurence, S., & Margolis, E. (1999). Concepts and cognitive science. In E. Margolis & S. Laurence (Eds.), *Concepts: Core readings*. Cambridge, MA: A Bradford Book, MIT Press.
- Levi, J. N. (1978). *The syntax and semantics of complex nominals*. New York: Academic Press.
- Libben, G. (1993) Are morphological structures computed during word recognition?. *Journal of Psycholinguistic Research*, 22(5), 535–544.
- Madden, J. D. (2013). *Mind, matter & nature: A Thomistic proposal for the philosophy of mind*. Washington, D.C.: The Catholic University of America Press.
- Mercier, D.-J. (1950a/1916). *A manual of modern scholastic philosophy (Vol. I)*. St. Louis, MO: B. Herder Book Co.
- Mercier, D.-J. (1950b/1916). *A manual of modern scholastic philosophy (Vol. II)*. St. Louis, MO: B. Herder Book Co.
- 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.
- Osherson, D. N., & Smith, E. E. (1981). On the adequacy of prototype theory as a theory of concepts. *Cognition*, 9, 35–58.
- Perler, D. (2001). Essentialism and direct realism: Some late medieval perspectives. *Topoi*, 19, 111–122.

- Pinheiro, J. C., & Bates, D. M. (2000). *Mixed-effects models in S and S-Plus*. New York, NY: Springer.
- Prasada, S., & Dillingham, E. M. (2009). Representation of principled connections: A window onto the formal aspect of common sense conception. *Cognitive Science*, 33, 401–448.
- Pustejovsky, J. (1995). *The generative Lexicon*. Cambridge, Mass.: MIT Press.
- Quené, H., & van den Bergh, H. (2004). On multi-level modeling of data from repeated measures designs: A tutorial. *Speech Communication*, 43, 103–121.
- Quené, H., & van den Bergh, H. (2008). Examples of mixed-effects modeling with crossed random effects and with binomial data. *Journal of Memory and Language*, 59, 413–425.
- Reynolds, P. L. (2001). Properties, causality, and epistemological optimism in Thomas Aquinas. *Recherches de Théologie et Philosophie Médiévales*, 68, 270–309.
- Smith, E. E., & Osherson, D. N. (1984). Conceptual combination with prototype concepts. *Cognitive Science*, 8, 337–361.
- Smith, E. E., Osherson, D. N., Rips, L. J., & Keane, M. (1988). Combining prototypes: A selective modification model. *Cognitive Science*, 12, 485–527.
- Spalding, T. L., & Gagné, C. L. (2013). Concepts in Aristotle and Aquinas: Implications for current theoretical approaches. *Journal of Theoretical and Philosophical Psychology*, 33, 71–89.
- Spalding, T. L., & Gagné, C. L. (2015). Property attribution in combined concepts. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 41(3), 693–707.
- Spalding, T. L., Gagné, C. L., Mullaly, A. C., & Ji, J. (2010). Relation-based interpretation of noun-noun phrases: A new theoretical approach. *Linguistische Berichte, Sonderheft*, 17, 283–315.
- Spalding, T. L., Stedman, J., Hancock, C., & Gagné, C. L. (2014). Intentionality and the Aristotelian-Thomistic view of concepts. *Journal of Mind and Behavior*, 35, 245–262.
- Wippel, J. F. (2000). *The metaphysical thought of Thomas Aquinas: From finite being to uncreated being*. Washington, D.C.: The Catholic University of America Press.
- Wisniewski, E. J. (1997). When concepts combine. *Psychonomic Bulletin & Review*, 4, 167–183.
- Wu, L., & Barsalou, L. W. (2009). Perceptual simulation in conceptual combination: Evidence from property generation. *Acta Psychologica*, 132, 173–189.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

