



Discussion

On prototypes as defaults (Comment on Connolly, Fodor, Gleitman and Gleitman, 2007) [☆]

Martin L. Jönsson ^a, James A. Hampton ^{b,*}

^a *Department of Philosophy, Lund University, Kungshuset, Lundagård, 22222 Lund, Sweden*

^b *Psychology Department, City University, Northampton Square, London EC1V 0HB, UK*

Received 13 November 2006; revised 26 February 2007; accepted 28 February 2007

1. Introduction

In an interesting contribution to research on conceptual combination, [Connolly, Fodor, Gleitman, and Gleitman \(2007\)](#) (CFGG) tested a hypothesis that they identified as a common assumption of prototype models of conceptual combination, ‘the crucial default to the stereotype prediction’ (DS). Defaulting to the stereotype consists in assuming, barring information to the contrary, that the prototype corresponding to an adjective noun combination (AN), fully inherits the properties of the prototype corresponding to the noun (N) of that combination. For instance, the prototype corresponding to *uncomfortable sofa* should share all the properties of the *sofa* prototype, except for those properties that have to do with comfort. In contradiction of this hypothesis, they reported an experiment in which people’s willingness to accept that a property is true of the members of some class tended to be lower when the concept was modified (*uncomfortable sofas have backrests*) than when it was unmodified (*sofas have backrests*). According to CFGG, since all prototype models of conceptual combination require DS to be true, the empirical demonstration that DS fails should be taken as important evidence against the correctness of such models.

[☆] The authors acknowledge the financial support of City University, London, and Erik and Gurli Hultengrens’s Foundation for Philosophy (Sweden). We wish to thank Pernilla Asp, John Barry, Herbert Clark, Zachary Estes, Daniel Heussen and Lena Wahlberg for fruitful discussions of the research and comments on the manuscript. The authors contributed equally to the research.

* Corresponding author.

E-mail address: hampton@city.ac.uk (J.A. Hampton).

Prior to reporting their experimental results, CFGG offered an argument to the effect that DS is in any case non-optimal and ‘invites indefinitely many bad bets’. Thus, according to CFGG, prototype theory embraces a misleading strategy, liable to lead people to make false inferences about the properties of complex concepts. Together, the empirical demonstration that people do not follow DS and the argument that DS is a bad strategy provide sufficient reason to favor some form of the ‘classical’ model over prototype models as the best approach to conceptual combination. The model they favor involves the assumption that (in the default case) modifiers and nouns correspond to concepts which are combined without being transformed into a structure having both of the original concepts as parts, so that the concept LITHUANIAN HAIR corresponds simply to a structure having the concepts LITHUANIAN and HAIR as parts. The classical combinatorial process involves no propagation of the prototypical/stereotypical properties of HAIR to LITHUANIAN HAIR although world knowledge may be evoked *post-combination* in order to make informed judgments about the likely properties of the complex concept.

Our paper aims to shed additional light on this topic, first by drawing out the implications of some recently published data that has explored and extended CFGG’s empirical result, and then by considering a number of important points in their paper. We disagree with them on several counts. First, we show that the results they report greatly exaggerate the extent to which subjects actually refrain from using the DS strategy. In fact when given the opportunity, most of the time people choose to say that a property is equally likely of a modified as of an unmodified noun, even when two atypical modifiers are used. Second, many of the prototype models that they claim to discredit do not embody the DS strategy at all, certainly not in the sense of *rigidly defaulting to the stereotype* regardless of the influence of knowledge-based inferences. Third, the results they report, contrary to undermining a central property of prototype approaches to conceptual combination, actually support a thesis that all these approaches do have in common – that the prototypes corresponding to complex expressions depend on the prototypes corresponding to their parts. Fourth, their argument that DS is an inefficient strategy is weak, and an alternative principled argument is offered for why in fact one should use this strategy. Our argument goes further in justifying why confidence in the inheritance of default attributes should be moderated by the typicality of the modifier – a point in their data that CFGG do not directly address.

We will argue for each of these points separately in the following sections. The final section argues that their results provide little reason to prefer the classical approach over the prototype approach to conceptual combination.

2. Prototypes and combination

2.1. DS is the most common strategy for interpreting modifier-noun phrases

In support of their claim that subjects do not follow the DS strategy, CFGG reported a study in which people judged a bare plural generic sentence asserting a

property of a class (*Ns are P*) to be more likely to be true than the same sentence when the noun was modified by an adjective or noun (*MNs are P*). So, for instance, people rated the sentence *Penguins live in cold climates* more likely to be true than the sentence *Solitary penguins live in cold climates*. Modifiers lowered sentence likelihood judgments regardless of whether a single typical modifier (e.g. *flightless penguins*), a single atypical modifier (e.g. *solitary penguins*) or two atypical modifiers (e.g. *solitary migrant penguins*) were used.

CFGG implicitly assumed that “likelihood of truth” ratings given for a sentence of the form ‘*Ns are P*’ correspond more or less directly to the weight assigned to the property *P* in the prototype corresponding to the *N*. This assumption is needed in order to bring the results they report to bear on prototype models of conceptual combination, since for the latter the weight of a property concerns the contribution that a property makes to determining typicality and degree of membership in a conceptual category rather than the degree to which people will judge it to be generically true. Prototype models of conceptual combination make no direct predictions about likelihood ratings as such. In fact the equation of property weight with likely truth may be questionable (for example Sloman, Love, & Ahn, 1998, show that feature weights tend to reflect causal centrality rather than statistical frequency). Even given the assumption that feature weights are strongly correlated with ratings of likelihood, our first claim is that the results CFGG report greatly exaggerate the extent to which people refrain from the DS strategy.

An experiment reported by Jönsson and Hampton (2007, Experiment 2) used the same materials as CFGG and obtained judgments of likelihood using a more direct procedure. CFGG obtained judgments of modified and unmodified versions of any one sentence from different groups of participants. Clearly, a more telling test of the DS hypothesis would be to present a respondent with the two relevant sentences and ask them to say whether they are equally likely, or, if not, which is the more likely. This procedure was therefore the one used in this experiment. Modified and unmodified sentences were judged equally likely on a majority of trials in all conditions (60% for double atypical modifier, 61% for single atypical modifier and 69% for the typical modifier). Put simply, the most common strategy employed in the study was DS. After having completed their judgments the participants were unexpectedly asked to go back and justify their choices where they had indicated that they thought that one of the sentences was more likely to be true. The most common reasons for preferring the unmodified sentence were either that it seemed more pragmatically sensible (an explanation already offered by CFGG for the effect of typical modifiers), or that there were reasons based on world knowledge for expecting the modified sentence to be less true. These two reasons accounted for 77% of meaningful justifications for why the atypical modifier was judged to have reduced likelihood and 97% of justifications for why the typical modifier did. When offering the pragmatic explanation, participants would often provide the additional comment that the two sentences were in fact equally likely – it was just that it made more sense to state the more general case. It is fair to conclude therefore that only for some 23% out of the 40% of sentence comparisons where a preference was expressed (or less than 1 in 10 cases overall) did participants judge that the modifier reduced the likelihood of the

generic sentence without offering some alternative explanation other than failure to default to stereotype.

It may be argued that many of these knowledge-based accounts used indirect reasoning. For example people judged that unripe peaches were less likely to be juicy than regular peaches. However, it should be noted that many models of prototype combination assume that this type of effect is essential to explaining the ways in which modified concepts have been shown to differ from their unmodified parents. Hence, the existence of these effects is completely in accordance with most available models of prototype combination.

In sum, Jönsson and Hampton (2007) did find a small proportion of answers that showed increased uncertainty about atypically modified sentences. So there was an effect over and above that due to knowledge based inferences and pragmatics. But it was substantially smaller than that reported by CFGG. An explanation is therefore needed not of why participants don't choose DS, but rather of why participants don't *always* use DS in the absence of pragmatic or knowledge-based reasons to the contrary. The general position taken by CFGG would argue that DS should never (or at least rarely) be used since it is a bad strategy. The important point we make here is that an explanation is also needed for why (for example) when two atypical modifiers were applied to the subject noun, 60% of the time the sentences were nonetheless judged equally likely to the unmodified sentence, and for another 25% of judgments there were either pragmatic or knowledge-based reasons offered for a reduced likelihood for the modified sentence. Participants apparently *do* default to stereotype, and they do it very frequently.

2.2. Prototype models do not generally predict that people will use the DS strategy

The extent to which models of prototype combination actually predict use of DS also bears closer examination. Should such models not make the prediction of DS, then obviously a finding that people do not use DS would not be evidence against such models. Even though (as we have argued above) DS may be the most common strategy in combining concepts, it is still important to examine the claim made by CFGG that available prototype models actually do predict DS, since they could argue that DS still fails in a significant number of cases. To examine the claim we will discuss in more detail two of the models that were targets of the criticism (CFGG, pp. 5–6).

Smith, Osherson, Rips, and Keane's (1988) *selective modification* model for simple adjective–noun combinations assumes that prototypes corresponding to nouns are sets of features organized into various dimensions.¹ Since the details are provided by CFGG, we will not repeat them here. We aim simply to point out that the model does in fact predict a general reduction in the weight of features when a modifier is added to a noun concept.

The model proposes that when a modifier is added to a noun (e.g. *brown + apple*), the weight of the relevant dimension (in this case COLOR) is increased relative to the

¹ The terminology of Smith et al. (1988) has been modified for comparability.

rest. As the weight of one dimension increases, inevitably the *relative* weight of other dimensions must decrease. Weights must be relative because they reflect the contribution of a feature to determining typicality (and only typicality in the case of Smith et al.'s model). So as one feature plays a stronger role, so the remaining features (as a whole) must necessarily play a weaker role. Hence the weight of the feature “is round” will be predicted to be lower for *brown apple* than it is for *apple*. This effect may be quite small, depending on how strong the weight of the modified dimension was to begin with, and how many other features there are in the prototype. The higher the initial weight of the modified dimension, then the smaller the effect of the modifier will be. No attempt has been made to produce a quantitative model of these effects. But the fact remains that the Smith et al. model clearly predicts that relative weights for all non-modified dimensions should decrease when a modifier is added to the noun concept, and this is consistent with the results reported by CFGG. Of course as we stressed above a further inferential step is required to deduce that differences in weights will correspond to differences in judged likelihood of the properties being true.

It is therefore incorrect to say that Smith et al.'s model necessarily predicts that people must default to the stereotype, or that the model is disproved by the finding that properties are judged less likely to be true of a modified concept.

Our second model, Hampton's (1987, 1988) *composite prototype* model proposes a mechanism for the combination of prototypes in conjunctive combinations of noun concepts (see also Costello & Keane, 2000; Thagard, 1997). In addition to predicting typicality effects, unlike the Smith et al. model the composite prototype model also applies to judgments of category membership in conjunctive concepts – and correctly predicts the occurrence of sizeable non-logical (and hence non-classical) deviations from strict intersection in people's judgments. The model has been tested primarily with relative clause conjunctions (e.g. “a sport that is also a game”), although Storms, Ruts, and Vandembroucke (1998) have demonstrated that other intersective phrases (adjective–noun or noun–noun combinations) show exactly the same phenomena.

The model proposes that noun concept conjunctions are formed in six steps:

- (1) a composite prototype is formed by the union of the features of the conjuncts,
- (2) all features with centrality so high that they are deemed necessary for either conjunct (e.g. fish have gills) will also be necessary for the conjunction,²
- (3) other features are assigned the average of their weights for each of the conjuncts, (a feature is given a weight of zero for a conjunct if it is not part of that prototype),
- (4) features with low resulting weights are eliminated,
- (5) a consistency checking procedure is run (informed by general knowledge), possibly resulting in the elimination and addition of further features in order to improve coherence,

² The separation of features into necessary and non-necessary at this step can be avoided if the averaging function is replaced by a more suitable continuous function that maps maximal weight for the conjunct onto a maximal weight for the conjunction.

- (6) examples of the conjunction may also be retrieved from memory, and features of these may be added.

It should be clear from the above that the composite prototype model doesn't entail DS. The weight of a feature for the combined concept is a function of its weight for each constituent. Hence, if the modifier does not embody a given feature of the noun, that feature will be inherited with half its original weight in the modified concept, due to the averaging function (unless the feature is considered necessary). A reduction in weight for properties that are unrelated to the modifier is therefore to be expected in this model also. As for the Smith et al. model, a further inferential step is required to link differences in weights to differences in judged likelihood of the properties being true. It would therefore be too strong to claim that the composite prototype model *predicts* that people do or do not default to the stereotype. Either way, CFGG's suggestion that DS is a necessary prediction of the model is incorrect.

We have elaborated two models in some detail. Two other models, Murphy's concept specialization model (Cohen & Murphy, 1984; Murphy, 1988) and Wisniewski's (1997) augmented schema model provide for a much wider range of possible conceptual combinations. Interpretation of compound noun phrases such as *apartment dog* or *shark lawyer* require a more complex set of processes than are relevant to the simple cases of adjectival modification considered by CFGG. Without going into detail, we note that both models emphasize the early use of background knowledge and local context in selecting the way in which individual concepts will interact to produce a complex concept. Unlike the first two models, there is no part of the models that explicitly proposes that feature weights would be reduced in a modified concept, but since there is equally no proposal that feature weights should be unaffected by modification, the claim that the models necessarily predict DS is unjustified. Most critically it is certainly not the case that these models *rigidly* default to the stereotype, i.e. retain features in the prototype corresponding to a complex whenever they are not *explicitly* mentioned in the modifier phrase and regardless of information that may be available from knowledge based inferences.³

2.3. The effect of modifiers provides support for a prototype framework

Regardless of whether or not specific models of prototype combination predict the effect that CFGG reported, one might wonder whether the effect has any bearing on

³ CFGG do repeatedly make this claim. For instance, "... a central prediction is that for any dimension x of prototype structure N that is not *explicitly affected* or *contradicted* by the requirements of combinations that N enters into, psychological measures that reflect the value of x should register no reliable differences across different combinatorial conditions. In this way, prototypes resemble the stereotypical members of their reference sets in the '*unmarked*' case" (pp. 6–7), "interesting prototype models have been developed that assume that all dimensions and weights of the simple conceptual representations are held unchanged under modification unless they figure in it *specifically*" (p.12), "the claim of DS that conceptual combination entails the inheritance of stereotypical default values for features that do not figure *explicitly* in the combination". (p.13) and "*unmentioned* values of the prototype representation are inert under combination". (p.14, our emphasis throughout).

the feasibility of the prototype approach to conceptual combination in *general*. And contrary to CFGG's claims we think that the existence of the effect actually strengthens this approach.

If the meaning of a modified noun phrase is constructed from the prototype of the noun, then one would expect that features which are considered more likely to be true of the noun would also be judged more likely to be true of the modified noun phrase. In other words, differences in feature strength observed in the noun prototype should be propagated to the prototypes corresponding to the modifier-noun constructions. This prediction was tested in a close replication of CFGG's experiment, reported in Jönsson and Hampton (2007, Experiment 1). A correlational analysis reported in this study showed good evidence that the relative strengths of features for the head noun concepts were inherited by both the atypically modified, and the double atypically modified noun concepts ($r(38) = .41$ and $.45$, $p < .001$). There was therefore evidence that the strength of a feature for the noun prototype was predictive of its strength for the modified noun phrase prototype, so that the prototype model was a promising framework in which to explain and understand the effect.

An additional discovery suggesting that the effect may be best treated within a prototype framework is reported by Jönsson and Hampton (2006) where it was shown that the effect persists even when the relevant sentences are *universally quantified*. So, in addition to judging it more likely that 'Strawberries have seeds' than that 'Lithuanian strawberries have seeds' (when a preference was expressed), people also considered it more likely that 'All strawberries have seeds' than that 'All Lithuanian strawberries have seeds' even though this constitutes a logical fallacy (which we termed the Inverse Conjunction Fallacy). This result suggests that just as in the case of Tversky and Kahneman's (1983) well known conjunction fallacy, people are using similarity to prototypical representations (i.e., representativeness) in order to judge the truth of propositions. People may believe that all purple apples are apples, but they may still be unwilling to conclude that something that is true of all apples is also true of all purple apples. The effect proved very robust, and generalized to other forms of universal quantification such as "Every single strawberry" and "100% of strawberries". For more details see Jönsson and Hampton (2006).

Now the classical model might account for the inverse conjunction fallacy by arguing that it results from some form of post-combination inference about the world. However, whereas the effect can be easily accounted for on the assumption that people are operating with prototypes, there is nothing in CFGG's model which would account for such a result in anything but a post hoc manner. In fact given that the classical account makes much of people's understanding of class inclusion relations (as seen, for instance, in CFGG's own account of their effect), it seems that the inverse conjunction fallacy, with its apparent disregard for class inclusion, runs counter to the classical view.

2.4. DS is not irrational or inefficient as a strategy

In Section 3 CFGG (pp. 8–9) argued that DS is in principle a "bad bet". As more modifiers are added to a concept (they claimed), the distance of the new

concept from the original gets greater and greater, so that carrying forward stereotypical information is likely to lead to increasing amounts of error. No formal argument was presented, but a rather misleading illustration was given in which the distribution of the modifier M and noun N are shown as strongly negatively associated (that is most M were not N and most N were not M).⁴ If the argument is really about “any” modifier however, one cannot know whether that modifier will be positively or negatively associated with the head noun class. If N and M are positively associated, then one’s degree of belief that MN will have the same properties as N should be greater, whereas if N and M are negatively associated, then it should be lower. The “in principle” argument about DS therefore depends on the likely distribution of the statistical association between arbitrary M and N categories.

We propose an alternative “in principle” argument *in favor* of DS as a strategy. Suppose that the modifier M picks out an arbitrary subset of the N class. Then it will be the case that the distribution of properties in N will provide the best estimate of the distribution of properties in MN (that would be to restate the fact that MN is an arbitrary subset of N). DS would then be the optimum strategy. If asked to assess the probability of any property P in the arbitrary sample of N defined by MN, the best estimate would have to be the probability of P in the sampled population N.

But what about the degree of confidence that one should have in the estimate? Suppose now that different modifiers pick out a greater or smaller (arbitrarily selected) proportion of the set N, while their possible association with any property P remains unknown. The greater the proportion of N that is contained in the subset MN, then the better the estimate of the frequency of P in MN will be (for example if MN represents 100% of N, the estimate will be perfect). The rational strategy would therefore be to predict that the frequency of P in MN is the same as that in N, but at the same time to reduce one’s confidence in the statement “MN are P” according to one’s estimate of the proportion of N that are MN. Since an atypical modifier (*Lithuanian strawberries*) is likely to select a smaller subclass of N than a typical modifier (*red strawberries*), it is then rational to reduce confidence in the truth of the statements in proportion to the typicality of the modifier. We do not claim that this is the reasoning behind people’s actual judgments; in fact we doubt that it is. But we present the argument as a counter to the notion that DS is in some sense an irrational strategy to follow, leading to “indefinitely many bad bets”. Quite the reverse is the case. However it is necessary to differentiate the fact that features are inherited from the parent concept by

⁴ The figure was particularly misleading in showing the modifier and noun with distributions on some dimension D, intended to be the property in question. This representation was misleading since in the case of MN phrases the modifier will normally have much larger scope than the noun on any particular dimension other than the one it itself modifies. Consider PURPLE APPLES. Apples have a distribution for size, but even if one could plot the distribution of the size of purple things in general, it certainly would not look like the curve in CFGG’s Fig. 2.

default from the need to qualify one’s confidence in the likelihood of the features still being true of the modified concept.⁵

3. The classical model versus the prototype model of conceptual combination

The model of conceptual combination that CFGG favored, the classical model, is one where “concepts remain inert under combination” (CFGG, p. 2), and “all you get from your concepts and combinatorics is output denoting relations among sets, properties, or individuals (depending on the ontology assumed)” (p. 4). In contrast to prototype theory, their model of conceptual combination does not presuppose that concepts have internal structure or that the result of combination is some transformation of the original concepts. Instead the result of combination is a simple structure having the original concepts as proper parts. Since the concepts lack internal structure, the classical model cannot supply explanations of typicality effects, or other features of everyday understanding directly, but has to be supplemented by “pragmatic-inferential processes that draw on general knowledge of the world”. It is therefore curious that their results show a strong effect of the typicality of the modifier (an effect replicated by Jönsson & Hampton, 2007, Experiment 1). If all that is available to the machinery of concept combination is the pair of concepts and the syntax for combining them, where does the effect of typicality come from?

The answer is that for CFGG, much of the interesting data concerning conceptual combination can only be explained by appeal to a second step involving “application of a further set of pragmatic-inferential processes that draw on general knowledge of the world” (CFGG, p.4) – a process that they later describe as “the ineffable all-purpose all-knowledge comprehension schema that we call ‘pragmatic inference’” (CFGG, footnote 9). Since they offer no account of this step, and little hope that such an account will be forthcoming, they don’t have an explanation of the effect under consideration. Note that DS being a bad strategy is not something which follows from the classical model (as described by CFGG) since the model contains no information regarding what counts as a ‘good instance’. So even *if* judgments of likely truth turn out to be good estimates of feature weights, and even if it could be shown that current models of prototype combination lack sufficient mechanisms to predict the extent to which people refrain from DS, prototype theory would still

⁵ CFGG cite Springer and Murphy (1992) who showed that emergent features that are true only for a modified noun phrase (‘boiled celery is soft’) were verified faster than features inherited from the noun (‘boiled celery is green’). CFGG argue therefore that combination occurs without involvement of the noun prototype, and that property information accrues once combination is complete. However, a more recent study (McElree, Murphy, & Ochoa, 2006) using the more rigorous method of speed-accuracy curves based on rapid forced responses showed that emergent phrasal properties are available significantly later than noun properties. The Springer and Murphy study allowed participants to make the decision in their own time, so that their results were probably influenced by pragmatic effects similar to those observed by CFGG (“boiledness” is not relevant to celery’s being green, but it is relevant to its being soft.)

seem to be a better place to look for an explanation of the result than the classical model. To be fair to CFGG's model it is not designed to provide explanations of typicality related phenomena but rather of the productivity and systematicity of thought, phenomena which are usually cited as the reasons for why a theory of conceptual combination has to be compositional.

What are we then to make of the arguments concerning the non-compositionality of prototypes? CFGG remark, as have Fodor (1998) and Fodor and Lepore (1996) before them, that prototypes are non-compositional, that is that the prototype corresponding to a complex expression is not fully determined by the prototypes of the parts of that expression and the way they have been combined. We completely agree with this. To derive the prototype corresponding to a complex expression, general knowledge is often (although not always) required, in addition to information regarding the prototypes for the individual parts of the expression and their mode of combination. However, being non-compositional in this sense need not be an explanatory shortcoming of prototype theory. Compositionality is considered important because it affords an explanation of productivity and systematicity. So if prototype theory can offer explanations of these phenomena without requiring compositionality (as it is normally understood), the argument against prototypes as concepts is weakened. (See Hampton & Jönsson, *in press*, for a more extended discussion.)

To account for productivity, that is how a system of representation can contain 'an infinite amount of syntactically and semantically distinct symbols' (Fodor & Lepore, 2002, p.1), it should be sufficient to provide recursive rules for combining whatever one identifies as concepts. All the prototype models considered here embody recursive rules. The composite prototype model for instance (described in this section), can be applied recursively to generate a concept such as "Recreations that are games, hobbies, sports and dangerous activities". Wisniewski's (1997) model can generate chains of concepts such as "Apartment dog", "Apartment dog coat", "Apartment dog coat cleaner appointment cancellation notification" and so forth. Systematicity refers to how a system of representation can contain 'such families of syntactically and semantically related but distinct expressions such as 'John loves Mary', 'Mary loves John', 'John loves John' and so forth' (Fodor & Lepore, 2002, p. 2). Again, we claim there is nothing that prevents prototype models of concept combination from incorporating systematicity. For example the composite prototype model would imply that anyone who can understand "Sports that are also Games" should be able to understand "Games that are also Sports", or "Sports that are also Sports" and so on.

We do not deny that the classical model of concept combination has a role to play in accounts of human cognition. It lies at the foundation of some of the finest intellectual achievements of our species such as mathematics and logic. However we would argue that it is mistaken to assume that the meanings that we construct when we combine concepts in everyday language can be captured by the classical model in anything but the most trivial of senses. To understand the complexity of the process will require a model that incorporates richer representations and a set of combinatorial rules that reflects that richness.

References

- Cohen, B., & Murphy, G. L. (1984). Models of concepts. *Cognitive Science*, 8, 27–58.
- Connolly, A. C., Fodor, J. A., Gleitman, L. R., & Gleitman, H. (2007). Why stereotypes don't even make good defaults. *Cognition*, 103, 1–22.
- Costello, F., & Keane, M. T. (2000). Efficient creativity: constraints on conceptual combination. *Cognitive Science*, 24, 299–349.
- Fodor, J. A. (1998). *Concepts: Where cognitive science went wrong*. Oxford: Clarendon Press.
- Fodor, J. A., & Lepore, E. (1996). The pet fish and the red herring: Why concepts still can't be prototypes. *Cognition*, 58, 243–276.
- Fodor, J. A., & Lepore, E. (2002). *The compositionality papers*. Oxford: Clarendon Press.
- Hampton, J. A. (1987). Inheritance of attributes in natural concept conjunctions. *Memory & Cognition*, 15, 55–71.
- Hampton, J. A. (1988). Overextension of conjunctive concepts: evidence for a unitary model of concept typicality and class inclusion. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 12–32.
- Hampton, J. A., & Jönsson, M. L. (in press). Typicality and compositionality: The logic of combining vague concepts. In W. Hinzen, M. Werning, & E. Machery (Eds.), *Handbook of Compositionality*. Oxford: Oxford University Press.
- 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. (2007). *The modifier effect in within-category induction: Default inheritance in complex noun phrases*. City University, London. (submitted for publication).
- McElree, B., Murphy, G. L., & Ochoa, T. (2006). Time-course of retrieving conceptual information: A speed-accuracy tradeoff study. *Psychonomic Bulletin & Review*, 13, 848–853.
- Murphy, G. L. (1988). Comprehending complex concepts. *Cognitive Science*, 12, 529–562.
- Sloman, S. A., Love, B. C., & Ahn, W. K. (1998). Feature centrality and conceptual coherence. *Cognitive Science*, 22, 189–228.
- Smith, E. E., Osherson, D. N., Rips, L. J., & Keane, M. (1988). Combining prototypes: a selective modification model. *Cognitive Science*, 12, 485–527.
- Springer, K., & Murphy, G. L. (1992). Feature availability in conceptual combination. *Psychological Science*, 3, 111–117.
- Storms, G., Ruts, W., & Vandenbroucke, A. (1998). Dominance, overextensions, and the conjunction effect in different syntactic phrasings of concept conjunctions. *European Journal of Cognitive Psychology*, 10, 337–372.
- Thagard, P. (1997). Coherent and creative conceptual combinations. In T. B. Ward, S. M. Smith, & J. Viad (Eds.), *Creative thought: An investigation of conceptual structures and processes* (pp. 129–141). Washington, D.C.: American Psychological Association.
- Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: the conjunction fallacy in probability judgment. *Psychological Review*, 90, 293–315.
- Wisniewski, E. J. (1997). When concepts combine. *Psychonomic Bulletin & Review*, 4, 167–183.