Nominal Prototypes and Compositionality

Abstract

One of the paradigms in the formal representations of meanings/concepts that lexical expressions (can be used to) express are prototype-theoretical representations. However, prototype-based accounts face a number of objections in relation to whether they can form part of an account of such meanings/concepts. A number of these objections revolve around problems with composing prototypes. For example, taking simple Boolean operations, prototypical features arguably emerge for complexes of prototypes, such as conjunctions, that are absent in the individual prototypes to be composed (the conjuncts), and negating prototypes has been argued not to yield prototypes. Restricting discussion to concrete nouns, I argue that even if we remain neutral on the exact place for prototypes in semantic theory, we should still require that prototype-structures compose. I then provide a probabilistic frame-theoretic account of nominal prototype structures. Combining such structures under Boolean operations not only follows from simple Bayesian calculus, but also accounts for when and why we should expect to find emergent features under conjunction, and why negated nominal prototypes seem to lack prototypical instances.

1 Introduction

The 'classical' view of the meanings of, *inter alia*, concrete nouns, is that such lexical items have necessary and sufficient conditions for (truthful) application. In the 1970s, an alternative to the classical paradigm began to emerge, a dominant form of which was prototype theory (see, amongst others, Rosch and Mervis, 1975; Rosch, 1978). On this alternative conception, words do not have an all-or-nothing criteria for application, but instead have a more graded interpretation. On the prototype theoretical view, categories have more or less prototypical members. One motivation for a non-classical approach is that it makes sense to ask how good/how central a case something is as a P. So, for example, a chair is a better case of *furniture* than a wall clock. Prototype theorists claimed that it did not make sense to ask or answer such questions if the classical view were right because all putative Ps would either satisfy the necessary or sufficient criteria, or they would not. All individuals in the denotation of P are first-class citizen on the classical view. The prototype approach allowed discrimination within a category.

A brief summary of some of the formal and informal foundations of prototype theory will be given in §2.1. However, there is an important terminalogical distinction which needs to be made at the outset. In the following, I shall distinguish between *prototypes* and *prototype structures.* Whereas a cat or a dog might be a good *prototype*/*prototypical case* for a pet, I will use *prototype structure* to refer to the formal representations used in prototype theory. These have been, for example, lists of typical features, and later attribute-value matrices, or frame-based representations (for an early suggestion for frame-based prototypes, see (Barsalou, 1992)). This distinction is of importance since in the literature 'prototype' is used equivocally, and, as shall be made clear later, some arguments which are sound with respect to prototypes, are not sound with respect to prototype structures.

There have been a number of controversies regarding the role prototypes and prototype structures should play within models of semantics and cognition. Principle amongst these are: (i) That typicality to a prototype and membership in a semantic class are arguably independent of one another; (ii) Whether prototypes/prototype structures are concepts, parts of concepts, or no part of concepts; (iii) how, if at all, prototypes compose.

In §§2.2–2.3 a brief overview of issues (i) and (ii) will be given. In §2.4, as well as looking at compositionality objections to prototype theory, we will consider a broad-based proposal for prototype structure combination made in (Prinz, 2002, 2012), the Retrieval Composition Analysis (RCA) model. In §2.5, I will argue that we should be able to give an account of how prototype structures compose regardless of whether we take prototype-structures to be meaning representations or not (and so regardless of where one stands on issues (i) and (ii)). In part, this argument will rest on the claim that non-compositional combinatorial mechanisms suggested by Prinz's RCA model cannot explain the systematic way in which the information composes under Boolean operations.

In response to this demand, in §3, I propose using, probabilistic Type Theory with Records (prob-TTR, (Cooper, 2012; Cooper *et al.*, 2014, 2013)), a probabilistic, compositional, frame-based formalism to model prototype-like structures. I set forward a broadly frame-based account of prototype structures interpreted in a Bayesian manner. Such structures, I will show, compose under Boolean operators, and can very naturally explain some of the challenges related to composing prototype structures. For example, in §4.1, I show how compositing probabilistic prototype structures under conjunction can give rise to emergent features, and in §4.2 I show how negations of prototype structures can lead to structures which would not predict strongly prototypical instances.

2 Prototypes and Problems with Prototypes

2.1 Background and Early Work

In the 1970s, a number of psychologists, lead by Rosch (Rosch and Mervis, 1975; Rosch, 1978) took inspiration from the Philosophical Investigations:

> Consider for example the proceedings that we call 'games'. I mean board-games, card-games, ball-games, Olympic games, and so on. What is common to them all?—Don't say: 'There must be something common, or they would not be called "games"'—but look and see whether there is anything common to all.—For if you look at them you will not see something that is common to all, but similarities, relationships, and a whole series of them at that. [...] And the result of this examination is: we see a complicated network of similarities overlapping and criss-crossing: sometimes overall similarities, sometimes similarities of detail. (Wittgenstein, 1953/2009, §66)

In (Rosch, 1978) Wittgenstein's remarks are explicitly linked with a discussion about boundaries and necessary and sufficient conditions. As Rosch sees it, to go in for necessary and sufficient conditions talk is to go in for boundaries talk. If, alternatively, concepts are seen as relating to a family of criss-crossing, overlapping features (sometimes of similarity and sometimes of detail), it was argued to be misguided to expect a concept to have a boundary (for to do so would just be to expect a set of necessary and sufficient conditions). Indeed, Rosch sees Wittgenstein's insight in being that one can judge how clear a case something is without knowing anything about boundaries. Following this line of thought, Rosch and others developed an alternative to the classical, necessary and sufficient conditions, accounts of semantic categories, one put in terms of *prototypes* and representations of structures which give rise to typicality judgements, what I call *prototype structures*.

Psychologists such as Rosch pointed out that, when asked, people make fairly predictable (proto)typicality judgements about what kinds of objects fall under some category, in that some types of objects are judged to be more typical of some superordinate category than others. It turns out, for example, that apples get judged as more typical than plums with respect to being fruit. Importantly, however, psychologists such as Rosch were not only interested in the objects/subcategories classified as typical, but were also interested in grounding these typicality data in some further theory (Rosch and Mervis, 1975). They therefore asked themselves what kinds of stories about representation, properties in the world, and learning would be compatible with the data they had collected about typicality judgements. Rosch and Mervis's (1975) proposal was to explain typicality judgements in terms of correlations perceived in ones environment. Apples, for example, are perceived to tend to have many properties or features shared with other fruits (edibility, approximate shape, tendencies toward colour). Plums, they found, do not tend to have quite as many shared features as apples (and presumably bananas even less so). So typicality qua superordinate category was suggested to be explained in terms of greater numbers of features shared with other members of the superordinate category. In this early work, no strong claims were made about what determines category membership. On the contrary, it is explicitly claimed in (Rosch and Mervis, 1975) that describing relative closeness to prototypes is an alternative to thinking of categories as 'logically bounded entities'.

What emerges from this view is a statistical approach to detailing semantic representations, or to be more neutral, a statistical approach to detailing representations of information exploited in communication and cognition. For prototype theorists such as Rosch, this, very broadly turns out to be a set of weighted features: how probable it is that something in one category will have some feature relative to how probable it is that something of a contrasting category will have it. It is this element of Rosch's work that I will adopt when it comes to modelling prototype structures for concrete nouns. Not, therefore the idea of a prototype *per se*, but models formed with stochastic *structures* that explain (proto)typicality judgements, and can themselves be grounded in a statistical account of how the relevant nouns are learnt.

A similar empirical approach is not only present in the early prototype theory literature, but may also be found in the approximately contemporary lexical frame-semantics approach ((Fillmore, 1975, 1976) amongst others). For example:

> Frame semantics comes out of traditions of empirical semantics rather than formal semantics. It is most akin to ethnographic semantics, the work of the anthropologist who moves into an alien culture and asks such questions as, 'What categories of experience are encoded by the members of this speech community through the linguistic choices that they make when they talk?' A frame semantics outlook is not (or is not necessarily) incompatible with work and results in formal semantics; but it differs importantly from formal semantics in emphasizing the continuities, rather than the discontinuities, between language and experience. (Fillmore, 1982, p.111)

The basic idea that we may take from both approaches is that if we can get a grasp on how the uses of words tally with our experience of the world, and what we can perceive in it, we might be able to explain some of the rich associations people make with respect to nominal categories.

As I shall argue, we can adopt the stochastic approach without having to decide whether such structures are, strictly speaking, models of nominal meanings/concepts, or not. However, a number of reasons for not identifying the two have been put forward. It is to some of these we now turn.

2.2 Typicality versus Membership

One early problem that arose for prototype-theoretic accounts was whether or not a measure of typicality is indicative of a measure of membership. This question has generally been raised in relation to vagueness by comparing the degrees of membership of a concept with degrees of typicality for a concept. One position is that the logical notion of concept membership is independent of the psychological notion of typicality (Osherson and Smith, 1982, 1997; Kamp and Partee, 1995). For example, (Osherson and Smith, 1982) argue that typicality is fundamentally epistemic, whereas membership is fundamentally logical. Somewhat similarly, (Rey, 1983) argues that prototypes best indicate *epistemic access* to metaphysically distinct extensions, but do not determine such extensions. However, others have argued for a closer connection between typicality and membership. For example, (Hampton, 2007) claims that measures for the two can be drawn from the same single measure of *similarity*.

2.3 Concepts or Not?

A similar, but distinct issue is whether or not prototypes or prototype structures are (adequate representations of) concepts, and whether certain concepts have prototypes (see, for example, Rey, 1983; Kamp and Partee, 1995). As I suggested in §2.1, at least some elements of Rosch's early work on prototypes did not directly identify prototypes with concepts, at least not if 'concepts' were meant to do everything that the classical view demanded of them. Rather, the goal was to explain how both prototype/typicality judgements, and seeming gradedness in concept membership arose (via a probabilistic account of learning and information processing). Nonetheless, overtime, emphasis within the psychology literature shifted and prototype structures started to be mooted *as* representations of concepts.

Most defenders of prototype theories now reject this view, however. Instead, they hold that prototypes/prototype structures are but one part of the story about the structures of concepts. For example, (Prinz, 2002) defends the view that prototypes are default concepts.[1] Even in earlier work, one finds introductory descriptions distancing the authors from a prototypes-as-concepts view:

> Research on natural concepts, such as *apple* and *fish*, has led to the conclusion that PART OF the mental representation of a concept consists of a 'prototype', roughly, a description of the best examples or central tendency of a concept (Smith *et al.*, 1988, p. 485) (my small caps)

There is not total contemporary consensus however. For example, Hampton ((Hampton, 1995, 2007) amongst others) defends the view that similarity to a prototype can account for typicality, vagueness, and concept membership, thereby forming a theory of at least some concepts. However, even Hampton does not endorse a strong form of the prototypes-are-concepts thesis:

> psychologists have never claimed that ALL concepts are represented in the mind as prototypes. Mathematical concepts such as i, the 'imaginary' root of -1, are clearly not represented as prototypes, and no-one has ever claimed that they are (Hampton and Jönsson, 2012, p. 398)

One of the reasons for doubting that concepts/meanings can be adequately represented with prototype structures is that prototypes and prototype structures do not always appear to be compositional. For the rest of this section, we will examine the details for such claims, and of a recently advocated model for prototype composition. Finally, I will present an argument for why there is a requirement for prototype structure compositionality irrespective of whether prototype structures are adequate (parts of) the representations of concepts/meanings, or of whether they determine category membership.

2.4 The Compositionality Problem

That representations of meanings are compositional is a standard assumption within the formal semantics literature and goes back at least as far as Frege. Kamp and Partee put the principle of compositionality succinctly:

> The meaning of a complex expression is a function of the meanings of its parts and of their syntactic mode of combination. (Kamp and Partee, 1995, p. 135)

Two main arguments for compositionality in natural language semantics are based on the idea that compositionality explains two other properties *productivity* and *systematicity*. From finite number of basic/lexical meanings, we seem able to express an infinite number of distinct complex meanings (productivity). Being able to grasp/express a thought seems, at least sometimes, to be a sufficient condition for being able to grasp/express others. For example, being able to grasp/express that aRb seems to be sufficient for being able to grasp/express bRa (systematicity). Compositionality, that the interpretations of all complex objects can be derived from the interpretations of some finite

[1] Empirical evidence against this view is presented in (Connolly *et al.*, 2007). See (Prinz, 2012) for a response.

number of simple objects plus some finite process of combination, is implied by the requirements for productivity and systematicity at least insofar as it is a parsimonious explanation for why languages should be both productive and systematic. I simply assume, for the purposes of this paper, that productivity, and systematicity are central properties of language and cognition. I will also not dispute the arguments from systematicity and productivity to compositionality.

The problem, for prototype-based accounts is that prototypes/prototype structures are, arguably, not compositional. The most forthright in this attack has been by Fodor and LePore (1991, 1996) and Fodor (1998, ch. 5). Non-compositonality is then used to separate concepts from prototypes. The core argument is:

(P1) Concepts must be compositional.

(P2) Stereotypes/prototypes are not compositional.

(C1) Concepts are not stereotypes/prototypes.

Justification for (P2) comes in different forms. I will consider two. The first is based on Boolean operators. If prototypes are compositional, then when linked with Boolean logical constants, the result should be a prototype. However, this seems not to be the case. For example, there might be a stereotype for *pet*, but there does not seem to be a stereotype for *not a pet*. Another argument is based on *emergent features*. Prototypes are defined, at least in part, in terms of features. However, it is not stereotypical for pets to be goldfish, and it is not stereotypical for fish to be goldfish, but *goldfish* is a highly stereotypical feature of *pet fish*, so it is hard to see how a *goldfish* can be a function of the features for *pet* and *fish*.

At this point, we need to take care with respect to the distinction between prototypes and prototype structures. The above arguments tend only to use the locution 'prototype' (see, for example, Fodor and LePore, 1996, p.260-3). However, the arguments for (P2) differ greatly when interpreted with respect to prototypes as opposed to prototype structures. Following in the tradition of (Rosch and Mervis, 1975), we have been concentrating on prototype structures (from which can be read or inferred typical instances/prototypes. Critically, it is no objection to this type of prototype theory to say that Boolean operations do not map prototypes to prototypes, provided that they map prototype *structures* to prototype *structures*. In the case of negation, one must also explain why the prototype structure that is the output to the function explains why one is left with no strong typicalities. In the case of emergent features, for example when two prototype structures are conjoined, one must also explain why the prototype structure that is the output to the function has strong typicalities which are not strong typicalities in the prototype structures that are inputs to the function.

There is a good deal of literature in response to the core argument, as well as to the arguments for (P2). In the first instance, as we have seen, some accept the conclusion of the core argument, but emphasise that prototypes/stereotypes are not meant to be identified with concepts (but only with parts of concepts) (Osherson and Smith, 1997), something that the argument does not disprove. Others, attack the basic argument (Robbins, 2002; Prinz, 2012). For example, (Prinz, 2002, ch.

11), (Prinz, 2012) argues that there are both mandatory and potential forms compositionality, that prototypes are potentially compositional, and that concepts need only be potentially compositional.

However, it has been suggested that Boolean operations on prototype structures can be explained non-compositionally. I will argue that this is not always the case, and, furthermore, that, at least for Boolean combinations, irrespective on the semantic status we assign them, we need to provide an account of how prototype structures compose.

2.4.1 The RCA model for prototype structure combination

Prinz (2002) suggests an RCA (Retrieval Composition Analysis) model. The reason why we will consider this model is that it adopts and combines together many elements of other proposals such as those in (Smith *et al.*, 1988; Hampton, 1991; Wisniewski, 1997). As such, it is a good summary of the means of combining prototype structures that have been suggested.

The RCA model predicts that prototype structures compose only given the failure of other means of concept combination available in the *Retreval* stage, in which, Prinz suggests, there are two retrieval strategies. On the one hand we can *search memory* for a previously stored compound concept. If not found, we can try *cross listing*, or a search for exemplars that fall under both prototypes. This, he claims covers cases such as *pet fish*, and thereby explains how emergent features, such as being a goldfish, arise. Only then, if still without success, do we combine prototype structures compositionally.

For the *Composition* stage, Prinz again suggests two two strategies: *alignment and integration* and *feature pooling*. Unlike traditional accounts of compositionality, both of these means of composition are context-sensitive. In this instance *context sensitivity* means that the contribution the meaning of a lexical item makes in the composition function may differ depending on the other inputs and/or the meaning it has in isolation. Prinz argues that this form of composition is still compatible with Fodor's systematicity requirement, and, as such, is an acceptable form of compositionality.[2]

Alignment and integration is based on the model in (Smith *et al.*, 1988). This *aligntegration* is meant to work for combinations such as *red fruit* as well as more artificial ones such as *squirrel snake*. In both cases, one aligns the shared features within the two prototype structures and integrates their values. For *red fruit*, this can be fairly straight forward. *Red* contributes some form of adjustment on the colour/red feature weighting already present in *fruit* (as detailed in (Smith *et al.*, 1988)). Prinz is less specific about how the process works for less standard cases such as *squirrel snake*:

> To align *squirrel* and *snake*, one looks for a dimension value in *snake* that is most similar to *squirrel*. Suppose that *snake* has the attribute *diet*, describing what snakes typically eat. This attribute may have a number of different values, including *mice*. Since squirrels have more in common with mice than other things represented in the *snake* concept, *squirrel* is aligned with the diet dimension and replaces *mice*. (Prinz, 2002, pp.303-4)

Even without full details, this form of context-sensitive 'composition' seems valuable. After all, *squirrel snake* could mean a squirrel-eating snake, but could just as easily be inferred to be, *inter*

[2]Since what I will suggest does not stand or fall on this argument, I will not assess this argument here.

alia a fake snake for deterring squirrels (some market stall holders place rubber snakes on their fruits to deter the seagulls, the same done for squirrels could make the rubber versions count as *squirrel snakes*). Whether we adopt a prototype-theoretic model or not, given this variability, we should not require the interpretation of *squirrel snake* to have a context-insensitive means of composition.

Although the focus of this paper is on nouns and $N \cdot N$ combinations, it is worth mentioning that a similar context-sensitivity point holds of many $A \cdot N$ combinations. Gradable adjectives (*red, tall, long* etc.) are almost universally accepted to be interpretable only in relation to context (an influential version of which in terms of thresholds and comparative classes has been developed by Kennedy (see, for example, Kennedy, 1997, 2007)). At the very least, context (either lexically or non-lexically provided) determines a comparison class relative to which the interpretation of the $A \cdot N$ phrase can be derived. In other words, the meanings of the lexical items and their mode of composition underdetermines the interpretation of the complex.[3]

The second composition strategy, *feature pooling*, which can be performed in parallel with aligntegration, is based on (Hampton, 1991). Feature pooling involves, in simple terms, combining features and dropping those with low weighting. For example, *house boat* retains some pooled features *place of domicile*, *floats on water*, but others are dropped *has a backyard*, *has sails*.

Prinz also suggests that feature pooling can account for *Hybridization* examples (Wisniewski, 1997). For Wisniewski, hybridization cuts across both Boolean and non-Boolean combinations such as the conjunction *musician painter*, and the non-conjunction *robin canary*. This means that, on the RCA model, sometimes Boolean combinations are derived via Retreval, and sometimes via a form of context-sensitive combination. On the account to be developed below, Boolean combinations can be accounted for with a single (non-context sensitive) composition function (one function for each Boolean connective). In the next section, I shall argue that, a fully compositional means of combining prototype structures under Boolean operations is needed whether or not prototype structures are taken to be models of meanings. In §3, I will show how this is formally realisable using fairly basic Bayesian calculus.

The final stage in the RCA model is *Analysis*. The result of context-sensitively combining prototype structures via aligntegration or feature pooling can leave gaps and incoherencies. One of Prinz's examples is that we might expect from *apartment dog* something both big and small. In the analysis stage, we are meant to bring background knowledge into play, and apply some pragmatic reasoning to settle this inconsistency.

2.5 Why the meaning question is irrelevant for the compositionality problem (Why we should anyway want prototype structures to compose)

Underpinning the compositionality problems raised for prototype theory is not the suspicion that prototype structures play no sort of cognitive role, or that they have no role in a theory of communication. Rather, the driving concern, for both compositionality and other objections, seems

[3] As has been pointed out, the N does not always provide the comparison class, hence the need for further context. This is demonstrated in examples such as, "Kyles car is an expensive BMW, though its not expensive for a BMW. In fact its the least expensive model they make." (Kennedy, 2007)

to have been that prototype structures are not themselves good representations of concepts/meanings. For example, (Rey, 1983) argues that prototypes and stereotypes tell us something about the epistemic access we have to concepts, the necessary and sufficient conditions for which may be external. We could, then, take such objections in earnest and see prototype structures merely as representations of the knowledge and beliefs we have about (the denotations of) concepts/meanings, but do not themselves constitute meanings/concepts.

The trouble is that such disputes over the question of what meanings *really* are, can end up in a, possibly unhelpful, dispute over the boundaries between the semantic and the epistemic/doxastic. One way of trying to avoid such disputes is return to some ideas presented in *The meaning of 'meaning'* (Putnam, 1975). The headline story, or at least what seems often to be taken as the headline story, in that article, is that *meanings ain't in the head*. This is a substantial philosophical claim, and it is surely right that it has been debated. However, the literal headline story of that paper is about the meaning of 'meaning'. Putnam's point about the various meanings of 'meaning', is that, even if there is a sense of 'meaning' on which to say what a word w means is to say what the true conditions of application for w are, there is (also) a sense of 'meaning' on which one can answer questions like 'What is the meaning of w?' and one's answer need not state the necessary and sufficient conditions for the application of w. For example, one could give a rough-and-ready description, or point to or describe an example of a w.[4]

Perhaps, despite Putnam's pleas to treat the word 'meaning' with care, nowadays 'meaning' is too loaded a term. However, his point still stands when put without using the m-word. If we want to get a good account of communication, we will have to take various factors into account. It is a matter of controversy whether there are such things as necessary and sufficient conditions for the application of nominal expressions.[5] Therefore there is dissent about whether or not such necessary and sufficient conditions are utilised by agents in communication, assertion, understanding etc. What is not controversial is whether other things that we know, believe, hold true, or assume, play a role in our communicative endeavours. Something about black and yellowish stripes and being a big cat plays a role in communicative exchanges using the word 'tiger' irrespective of whether the having of such stripes, or the being a big cat is necessary or sufficient for being a tiger. Whether or not we call such defeasible assumptions about tigers '(part of) the meaning of "tiger"' is to some extent by-the-by if we want to understand the rich variety of communicative possibilities available using the word 'tiger'. Such information plays a role in communication whether we call it 'meaning' or not and one proposal for representing this other information is via prototype structures. If the prototypes proposal does not suffer from other independent problems, therefore, such structures are good candidates for playing a role in our formal representations of information flow and utilisation in communication whether we call such structures "representations of meanings" or not.

However, once we accept this, the need for compositionality, in some form or another, does not

[4]A similar point was made earlier by Austin (1970). However, for Austin, once one has exhausted a story about how one might ordinarily reply to questions like "What is the meaning of w?", there is little sense in then asking, "But then what is the-meaning-of-w?".

[5]For assent, see (Lewis, 1970; Cappelen and Lepore, 2005; Borg, 2004) amongst many others. For varying levels of dissent see, amongst others, (Travis, 1978; Recanati, 2004; Rayo, 2011).

go away. Say that for some words $w_1, w_2, ..w_n$ we have some clusters of information $I_{w_1}, I_{w_2}, ..., I_{w_n}$ which individually may or may not be called the meanings of $w_1, w_2, ..w_n$ respectively. Then we take our favourite syntax-semantics interface which takes as input some concatenation of words, and if the words can sensibly be concatenated, it outputs, minimally, some logical form built out of logical and non-logical constants (and maybe some variables). In some way, then, our clusters of information $I_{w_1}, I_{w_2}, ..., I_{w_n}$ are to be associated with the non-logical constants in this output, regardless of whether they are the meanings. Now, any reasons we might have for wanting to be able to productively and systematically combine whatever *really are* the meanings of $w_1, w_2, ..w_n$, are also reasons for wanting to productively and systematically combine the clusters of information $I_{w_1}, I_{w_2}, ..., I_{w_n}$. My argument for this is a disjunction elimination:

(P1) Either $I_{w_1}, I_{w_2}, ..., I_{w_n}$ are the meanings of $w_1, w_2, ..w_n$ or they are not.

(P2) If they are, then $I_{w_1}, I_{w_2}, ..., I_{w_n}$ should be compositional.

(P3) If they are not, then $I_{w_1}, I_{w_2}, ..., I_{w_n}$ should be compositional.

(C1) $I_{w_1}, I_{w_2}, ..., I_{w_n}$ should be compositional.

The crucial premises are (P2) and (P3). I take them in turn.

For (P2), I assume, in general, that arguments for the compositionality of meanings are sound. This is not uncontroversial, but it is not in the scope of this paper to doubt the compositionality of semantic entities. Hence, (P2) follows from this assumption.

To see why (P3) holds, let us consider a situation where it did not. In such a situation, one would have the non-meaning bits of information nonetheless linked to the words and the meanings they have. For complex expressions, whatever the meanings of the words in the expressions were, they would compose via a finite set of rules to form a complex meaning. However, denied any means of composing, $I_{w_1}, I_{w_2}, ..., I_{w_n}$, would simply float, unconnected to this complex in any meaningful way, and any connections or overlaps in the information would not be exploited (beyond, perhaps, simple feature cross listing of exemplars). For example, if part of the (non-meaning) information associated with 'tiger' is *big stripy feline*, part of the (non-meaning) information associated with 'wolf' is *dog-like pack animal* and part of the (non-meaning) information associated with 'cub' is *individual smaller and younger than an adult*, then we should not expect there to a systematic (non-meaning) output for complexes such as 'tiger cub' and 'wolf cub' such as *individual smaller and younger than an adult big stripy feline* and *individual smaller and younger than an adult dog-like pack animal* respectively. Yet we do frequently see such systematic, productive combinations. Equally, if some of the (non-meaning, defeasible) information associated with both 'singer' and 'songwriter' is *musically talented*, without any link between this information and some manner of composition, there would be no particular reason to associate the information *musically talented* with the locution 'singer songwriter'.

If we assume that such information is not part of the meaning of the associated lexical item, it is therefore a requirement on an adequate theory of communication that we can nonetheless explain how we arrive at the (non-meaning) information associated with complexes, albeit possibly in parallel with an account of semantic compositionality.

To narrow any remaining room for doubt, we can consider whether the result of prototype structure combinations could be gained non-compositionally. There are three non-compositional processes described in (Prinz, 2002). Compound search, exemplar cross listing, and analysis. we might also add to this list free pragmatic reasoning. Now take the following example. A learner acquires the, albeit defeasible, information that tigers are stripy, that wolfs are doglike, and that cubs are smaller than adults, then she is presented with the novel compound 'wolf cub'. Compound search is unavailable, because the compound is novel. Cross-listing is also unavailable, because even if she has an exemplar for a tiger (an adult tiger), an exemplar for a cub (a tiger cub), and an exemplar for a wolf (an adult wolf), there is no shared exemplar for both *wolf* and *cub*. All that is left is *analysis* and free pragmatic enrichment. Analysis won't do, because it is meant to be applied only to incoherent outcomes of composition. Free pragmatic inference might account for the result of the combination, however, to assume that free pragmatic inference accounts for the interpretation of all noun combinations would have a high cost. First it might endanger the systematicity of such processes. For example, if I_3 is the 'correct' information to associate with $N_1 \cdot N_2$, and I_4 is the 'correct' information to associate with $N_2 \cdot N_1$, the ability of an agent to reason to I_3, thereby interpreting $N_1 \cdot N_2$, does not, on its own, imply she or he would be able to reason to I_4, thereby interpreting $N_2 \cdot N_1$. Second, over reliance on free reasoning runs the risk of making the learning of compounds implausible, since the number of pragmatic inferences available (given access to sufficient quantities of extra background knowledge), would, in principle, be unbounded.

This leaves two possibilities. Either combinations of prototype structures combine via context-sensitive mechanism such as aligntegration or feature pooling, or they combine in a way determined by the input prototype structure(s) and some function on that structure provided by the mode of combination.

This may be enough to secure (P3) depending on how we interpret 'compositional'. If we follow (Prinz, 2002, 2012; Hampton and Jönsson, 2012) in allowing that compositionality can be context-sensitive, then (P3) has been established by ruling out other combinatorial mechanisms. If we wish to retain Fodorian strictness with respect to context-insensitivity, (P3) is harder to establish, since it would mean also ruling out aligntegration and feature pooling and I do not have any decisive way of ruling such mechanisms out.

However, we do have some weak evidence for favouring a context-insensitive means of combination for *Boolean* complexes. There are cases of $N \cdot N$ combinations, described above, such as *squirrel snake*, where numerous interpretations are available, and so the generation of the resultant prototype structures may in these cases be derived from more than just the input prototype structures and their means of combination (context-sensitive composition). Yet *squirrel snake* is a good example of a *non*-Boolean compound (squirrel snakes are not snakes and squirrels). Indeed, we seem to get systematic results from nominal prototype structure combinations exactly when the logical form of what is said includes Boolean operations. This can most clearly be seen with negation.

As a general rule, negating a prototype structure eliminates the typicalities associated with the non-negated prototype structure. Conjunction is more complex, since conjoining prototype structures can sometimes, but not always give rise to highly typical features that were not highly typical of either conjoined prototype structure. Nonetheless, even for conjunction, there seem to be a few rough rules of thumb. First, shared features are retained under conjunction (*musically talented* for 'singer songwriter'). Second, if one conjunct is neutral with respect to a feature that the other conjunct strongly predicts, then the conjunction inherits the stronger feature (*dog-likeness* for 'wolf cub'; *big stripy cattiness* for 'tiger cub). Finally, in cases where features are in conflict, we tend to see emergence (*species type* for 'pet fish').

These observations are only very informally articulated, but I will show how the emergence or non-emergence of features can be shown to be entirely systematic and productive when interpreted via a Bayesian model. If this is possible, then, although weak, it would also support the conclusion that prototype structures are (strictly) compositional under Boolean operations, no least because a strictly combinatorial account is more parsimonious than appealing to a combination of compound searching, cross-listing, aligntegration, and feature pooling for complexes with the same logical form.

We thereby have our conclusion, whether meanings or not, there is good reason to want and expect prototype structures to compose (in some sense of 'compose'), because we do retrieve the kind of information encoded by complexes, even though this information is not plausibly derivable from non-compositional mechanisms alone. Furthermore, there are definite patterns in the effects of applying Boolean negation to prototype structures, and arguably also (complex) patterns that arise from conjunctions, which weakly support a strictly compositional interpretation for combining prototype structures with these Boolean logical constants. In the remains of this article, I put forward a Bayesian model of prototype structures, and will provide a strictly compositional method for applying Boolean connectives and operators to prototype structures.

3 Probabilistic frames as prototype structures

The basis of the proposal for prototype structures to be made here will be inspired by frame semantics (Fillmore, 1975, 1976) and frames as attribute value matrices (Barsalou, 1992). Early prototype structures were simply based upon feature lists. For example, *rose* might have *red*, *thorny* etc. However, now they are more often thought of as frames. Frames have attribute-value structure, rather than mere feature listing. So *rose* might have attributes *colour*, *texture*, which in turn have values such as *red*, *yellow*, *white*, and, *thorny* respectively.[6]

The simple idea here will be to interpret attribute-value structures in fairly straight-forward Bayesian terms, such that prototype structures will be describable as collections of conditional probability distributions and collections of prior probabilities. Figure 1., helps to show how such a prototype structure might look.

[6] The actual structure would be more complex. For example, different attribute-value pairings might be attached to superordinal categories such as *flower* and *stem*.

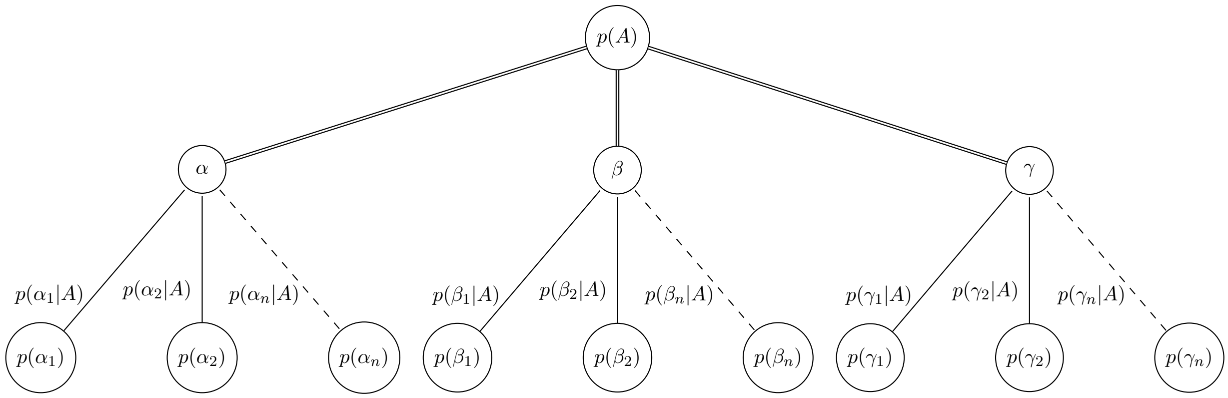

Figure 1: Probabilistic prototype frame for A, with attributes α, β, γ

The prototype structure in Figure 1 is for A (which in our case will be the label for the prototype structure associated with a noun). One element in this structure is a prior probability for the head category, $p(A)$ (I will say more about the source for such priors below). Double-lined connections link $p(A)$ to attributes α, β, γ. I assume that instances of the prototype must have a value for each attribute. For example, any prototype for *pet* must have a value for the attribute *(animal) species* (if we accept that pets must be some species of animal or another). In this sense, attributes can be seen as variables for the values they take. Single lines link values to attributes and indicate a conditional probability. In the example just described this might include *cat*, *dog* etc., but also other species such as *giraffe* or *haddock*. Each of these species will receive a prior probability, but crucially, I assume a dependency between each value and the prototype head. The values for each attribute will be assigned a probability conditional on the prototype category. These conditional probabilities will form a distribution. Therefore, for example:

$$p(\alpha_1|A) + p(\alpha_2|A)p(\alpha_1|A) + ... + p(\alpha_n|A) = 1$$

In sum, a probabilistic prototype structure for a category A will consist of three collections of probabilities:

1. a prior for A ($p(A)$);

2. priors for all values the attributes an A might take ($p(\alpha_1), ..., p(\alpha_n), p(\beta_1), ...p(\beta_n), p(\gamma_1)...$);

3. conditional probability distributions linking the prototype category to it's possible values (one distribution per attribute, e.g., $p(\alpha_i|A), ..., p(\alpha_n|A), p(\beta_i|A), ..., p(\beta_n|A), ...p(\gamma_1|A), ...$).

3.1 A very brief introduction to TTR and prob-TTR

In terms of a formal representation, I adopt a probabilistic variant of Type Theory with Records (TTR, (Cooper, 2012, 2010) amongst others), namely, prob-TTR (Cooper *et al.*, 2014, 2013). TTR is fairly unique in that it directly unites two often seemingly disparate fields within semantics, namely lexical frame semantics (Fillmore, 1975, 1976), and compositional semantics in the tradition of Montague (see e.g., Thomason, 1974). Furthermore, TTR and prob-TTR record types are designed to be able to represent semantic information, as well as information of a more doxastic and epistemic nature (see (Dobnik *et al.*, 2013) for an incorporation of TTR frames with perceptual information). Taking all of these factors together, given our deliberately neutral stance on the semantic or non-semantic status of prototype structures, this makes prob-TTR is the perfect tool to adopt to represent probabilistic, Bayesian prototype frame structures.

Records, in TTR, are approximate cognates for situations in Situation Theory. *Record Types* can be understood as *frames* in the sense of Fillmore (1976), but are also approximate cognates for *Situation Types/Infons.* On standard TTR (Cooper, 2012), a nominal predicate such as *pet* can, simplifying a lot,[7] be represented as in (3.1), which is a function from records of type $[x : Ind]$, to records of type $[c_1 : pet(r.x)]$. The $r.x$ characterises a form of binding, that the value for x in $pet(x)$ be the same as the value for x in $x : Ind$.

$$\lambda r : \left[\text{ x } : Ind \right] . \left(\left[\text{ c}_1 \;:\; \text{pet(r.x)} \right] \right) \tag{3.1}$$

After applying a record to the function in (3.1) will yield a record type (or a type-theoretic proposition). The following three records could be supplied as arguments to the function in (3.1) if the individuals a, b, c are typed as individuals:

$$r_1 = \begin{bmatrix} x & = & a \\ c_1 & = & p_1 \end{bmatrix} \;,\; r_2 = \begin{bmatrix} x & = & b \\ c_1 & = & p_2 \end{bmatrix} \;,\; r_3 = \begin{bmatrix} x & = & c \\ c_1 & = & p_3 \end{bmatrix} \tag{3.2}$$

The values for c_1 in the above records, p_1, p_2, and p_3, are *proofs*. For example, a proof of $pet(b)$ can be understood in situation-theoretic terms as a situation in which the individual b is a pet.

The result of applying the records in (3.2) to the function in (3.1) would be the record types/propositions in (3.3):

$$\left[\; c_1 \;:\; pet(a) \; \right] \;,\; \left[\; c_1 \;:\; pet(b) \; \right] \;,\; \left[\; c_1 \;:\; pet(c) \; \right] \tag{3.3}$$

At the heart of TTR is the notion of a type judgement. This is loosely based on the idea in (Austin, 1950/1979) that truth-conditions can be specified in terms of situations being of the appropriate type. In TTR, type judgements are of the form $r : T$ (a record r is judged to be of type T). Assuming that p_1, p_2, and p_3 are all proofs of $pet(a)$, $pet(b)$, and $pet(c)$, respectively, the following Austinian type judgements would be true:

$$r_1 : \left[\; c_1 \;:\; pet(a) \; \right] \;,\; r_2 : \left[\; c_1 \;:\; pet(b) \; \right] \;,\; r_3 : \left[\; c_1 \;:\; pet(c) \; \right] \tag{3.4}$$

In the probabilistic variant of TTR, prob-TTR (Cooper *et al.*, 2014, 2013) type-judgements can be linked to Bayesian semantic learning models. Our semantic learning data is insufficient to provide

[7] Amongst other things, the record type should include information relating to time and location.

us with determinate type judgements. Instead, we can only be more or less certain that an object is of some type (given our learning data). The (determinate) Austinian proposition $r : T$ is therefore replaced with a probabilistic equivalent, the probability that a record is of a type: $p(r : T) = k$ where $k \in [0, 1]$. This probability is meant to reflect how certain an idealized agent should/would be that something is of a type.

The functional nature of TTR is preserved within prob-TTR. Hence nominal predicates are still functions from records, to record types, except now, record types (propositions) are interpreted probabilistically, namely the probability that a record r is of the resultant type, namely:

$$p(r : \begin{bmatrix} x & : & Ind \\ c_A & : & A(x) \end{bmatrix}) = k \tag{3.5}$$

The exact value for this judgement may depend on various other factors. For example, call the above type T_1, if the record r is also of some other type T_2, and if a dependency exists between types T_1 and T_2, then part of the judgement in (3.5) will turn on the probability that a record r is of type T_1, given that it is of type T_2, namely, $p(r : T_1 | r : T_2)$.

4 Composing Probabilistic Prototype Structures

4.1 Conjunctions of Prototype Structures

In TTR, $N \cdot N$ conjunctions can be represented as functions from records that are witnesses of individuals to record types dependent on the argument record:

$$\lambda r : [\ x \ : \ Ind\ ]. \begin{bmatrix} c_A & : & A(r.x) \\ c_B & : & B(r.x) \end{bmatrix} \tag{4.6}$$

If the individual witness of the record provided as argument for the above function is a, the result is a record type in (4.7), the type of record/situation in which a is of the predicate type (p-type) A and the p-type B:

$$\begin{bmatrix} c_A & : & A(a) \\ c_B & : & B(a) \end{bmatrix} \tag{4.7}$$

The prob-TTR probabilistic type judgements for such a proposition are calculated in line with the standard Bayesian conjunction rule:

$$\begin{aligned} p(a : T \wedge T') &= p(a : T) \times p(a : T' | a : T) \\ &= p(a : T') \times p(a : T | a : T') \end{aligned} \tag{4.8}$$

Conditional probabilities are calculated in accordance with a type theoretic version of Bayes' Rule:

$$p(a : T|a : T') = \frac{p(a : T \wedge T')}{p(a : T')} \tag{4.9}$$

Of more interest to us however, is how information associated with p-types A and B in (4.7) might be combined. Importing the probabilisitic frame structure schematised in Figure 1, I assume, for each such p-type, a probabilistic frame prototype which encodes dependencies between the p-type and the values of its attributes, as well as the priors of the values. For a conjunction of types A, and B, we can simply apply Bayes' rule (a few times) to derive the probability of some attribute value, given this conjunction. This will be possible for any value as long as A and B share that value's attribute. Combinations can be provided for all values of attributes which A and B share. As we will see, this allows for values which are improbable given A and improbable given B, to be highly probable given A and B. In other words, we shall see how simple Bayesian calculations can yield emergent features.

Assuming a prototype structure for A and B, with some value α_i we can be shown to have all we need to calculate the probability of the value, given the conjunction. Let T_A, T_B, and T_{α_i} be the record types relating to the p-types A, B and α_i respectively. Then:

$$p(a : T_{\alpha_i}|a : T_A \wedge T_B) = \frac{p(a : T_{\alpha_i} \wedge T_A \wedge T_B)}{p(a : T_A \wedge T_B)} \tag{4.10}$$

There are, in (4.10) some values not directly provided by either the prototype structure for A or the prototype structure for B. However, using some simple Bayesian manipulation, these can be shown to be derivable from the values of these structures.

Although we need not assume that T_A and T_B are independent, we must make some form of independence assumptions. Here, I assume that $p(a : T_A|a : T_{\alpha_i})$ and $p(a : T_B|a : T_{\alpha_i})$ are independent. To motivate this assumption, take an example. A shared attribute of the type *pet* and the type $fish$ might be *species*. A value for *species* could be the type $goldfish$. Our independence assumption means that the probability that something is a pet, given it is a goldfish will not be affected by the probability that something is a fish given that it is a goldfish. This is surely right, since the former is wholly contingent (for example, on people's pet-keeping habits), whereas the latter is, if not necessary, somewhat hard to doubt (all goldfish are fish). With this independence assumption, we have:

$$p(a : T_{\alpha_i}|a : T_A \wedge T_B) = \frac{p(a : T_{\alpha_i}) \times p(a : T_A|a : T_{\alpha_i}) \times p(a : T_B|a : T_{\alpha_i})}{p(a : T_A \wedge T_B)} \tag{4.11}$$

Where $p(a : T_A|a : T_{\alpha_i})$ and $p(a : T_B|a : T_{\alpha_i})$ can be derived via Bayes' Rule from the prototype structures for A and B. Finally, $p(a : T_A \wedge T_B)$ can be calculated too. Since distributions sum to 1, we know that:

$$p(a : T_{\alpha_1}|a : T_A \wedge T_B) + p(a : T_{\alpha_1}|a : T_A \wedge T_B) + ...p(a : T_{\alpha_n}|a : T_A \wedge T_B) = 1 \tag{4.12}$$

Hence we can replace $p(a : T_A \wedge T_B)$ with a sum over all values in the attribute:

$$p(a : T_{\alpha_i} | a : T_A \wedge T_B) \quad = \quad \frac{p(a : T_{\alpha_i}) \times p(a : T_A | a : T_{\alpha_i}) \times p(a : T_B | a : T_{\alpha_i})}{\sum_n p(a : T_{\alpha_n}) \times p(a : T_A | a : T_{\alpha_n}) \times p(a : T_B | a : T_{\alpha_n})} \tag{4.13}$$

4.1.1 Example: pet fish

A plausible shared attribute for the prototype structures for *pet* and *fish* is *species*. The *species* attribute will have many values, but for simplicity, I assume there to be only five, *cat*, *dog*, *goldfish*, *haddock*, *tuna*. The relevant priors and conditional probabilities for these types would, in an implemented model, be linked to an agent's semantic learning data. In this toy example, I have chosen priors and conditional probabilities that reflect the following intuitions (which could be adjusted): cats and dogs can't be fish; goldfish, haddock and tuna must be fish; cats and dogs are far more probable instances of pets than goldfish; goldfish are far more probable instances of pets than haddock or tuna; haddock and tuna are far more probable instances of fish than goldfish.

These intuitions are reflected in the priors and conditional probabilities shown in Table 1 (the record types for these types will be represented as: T_{pet}, T_{fish}, T_{cat}, T_{dog}, $T_{goldfish}$, $T_{haddock}$, T_{tuna}). The final column in Table 1 shows the result of composing *pet* and *fish* with respect to the shared attribute *species*.

i	$p(a{:}T_i)$	$p(a{:}T_i \mid a{:}T_{pet})$	$p(a{:}T_i \mid a{:}T_{fish})$	$p(a{:}T_{pet} \mid a{:}T_i)$	$p(a{:}T_{fish} \mid a{:}T_i)$	$p(a{:}T_i \mid a{:}T_{pet} \wedge T_{fish})$
cat	0.01	0.4	0	0.8	0	0
dog	0.01	0.4	0	0.8	0	0
goldfish	0.005	0.198	0.2	0.8	1	0.91
haddock	0.01	0.001	0.4	0.02	1	0.045
tuna	0.01	0.001	0.4	0.02	1	0.045
pet	0.02					
fish	0.025					

Table 1: Composing *pet* and *fish* with respect to the attribute *species*

What Table 1 helps to show is that, from fairly innocent assumptions about what makes a plausible/typical pet and a plausible/typical fish, as well as what makes an implausible/a-typical pet, and an implausible/a-typical fish, we get an emergent feature, *goldfish.* Though fairly a-typical pets and fairly a-typical fish, goldfish can emerge as incredibly typical pet fish. Furthermore, this can be modelled as a direct result of composing two prototype structures under Boolean conjunction, interpreted according to classical (Bayesian) probability calculus.

Still to be elaborated is how some of the other 'rules of thumb' for prototype structure conjunction can be respected (see §2.5). Limitations of space prevent the provision of such details, however it should be fairly clear how, for example, feature strengthening (*musically talented* for *singer songwriter*), and feature adoption (*doglike* for *wolf cub*), will result from the Bayesian conjunction rule detailed above. I leave a demonstration of this and other examples for future research.

4.2 Negation of Prototype Structures

Recall that one of the compositionality-based objections was that Boolean operations on prototypes tend not to yield prototypes. For example, there might be a prototypical fish, but there is no, or at least no strong, prototype for *not a fish*. Again, we will not compose prototypes, but prototype structures. It will be perfectly in line with compositionality if the result of applying negation to a prototype structure yields something that does not have strong prototypical instances. All that matters is that the output of such a function is itself a prototype structure.

Hence, negations such as *not a pet*, *not a fish* will be functions from prototype structures to prototype structures, calculated on wholly Bayesian lines. Basic Bayesian negation holds in prob-TTR. In notation: $p(a : \neg T) = 1 - p(a : T)$. However, similarly to conjunction, where we derived the probabilities of attribute values conditional on the conjunction of predicates, for negation, we will derive the probabilities of attribute values, given the negated predicate type:

$$p(a : T_{\alpha_i} | a : \neg T_A) = \frac{p(a : T_{\alpha_i}) \times (1 - p(a : T_A | a : T_{\alpha_i}))}{1 - p(a : T_A)} \tag{4.14}$$

4.2.1 Examples: not a pet, not a fish

By way of example, we can generate prototype structures for $\neg pet$ and $\neg fish$ with respect to the attribute *species*. I re-use priors displayed in Table 1, the results are shown in Table 2.

Negating *pet* has numerous effects. First, it lowers values for common and somewhat common pets (cats, dogs, and goldfish) with respect to their priors. So, *not a pet* should make one expect a dog, cat or goldfish less than one would in a neutral context. It also leaves values for haddock and tuna unaltered. So, *not a pet* should make one expect a haddock or a tuna to the same extent one would in a neutral context (i.e. compared with ones prior probabilities).

Negating $fish$ also has numerous effects. First, it also leaves values for cat and dog unaltered.[8] So, *not a fish* should make one expect a dog or a cat to the same extent one would in a neutral context (i.e. compared with ones prior probabilities). Second, as one might expect, our examples of fish (haddock, tuna, goldfish) are ruled out completely by *not a fish*.

i	$p(a : T_i)$	$p(a : T_i \| a : \neg T_{pet})$	$p(a : T_i \| a : \neg T_{fish})$
cat	0.01	0.002	0.01
dog	0.01	0.002	0.01
$goldfish$	0.005	0.001	0
$haddock$	0.01	0.01	0
$tuna$	0.01	0.01	0
pet	0.02		
$fish$	0.025		

Table 2: Negating *pet* and *fish* with respect to the attribute *species*

Hence, we can capture the intuition that negating prototypes leaves one without typical instances. Aside from that which is excluded (fish species are ruled out by *not a fish*), for many attribute values, one may be left effectively with the same expectations as one's priors. What this represents is that

[8] This is not entirely accurate, to more decimal places, the new values are 0.00103.

there will, at least for many attribute values, be no strongly typical instances derivable from the negated prototype structure.

One potential problem with this model is that it does not seem to capture some remarks made by Prinz (Prinz, 2002) on negated prototypes. He cites evidence that, transposing for the above example, haddock and tuna should be fairly typical non-pets, but that something like an egg whisk should be a more typical non-pet. One possibility to accommodate such data is to make attributes probabilistic as well as their values (thereby leaving, for example, a slim probability that, say, a pet could be a non-animal/no species). However, I leave detailing this suggestion for further research.

5 Summary and Conclusion

The kind of probabilistic prototype structures detailed above were meant to capture some of the early insights of Rosch and Mervis (1975). There can be value in describing, in stochastical terms, interrelations between typical and atypical features of entities within superordinate categories, since the information represented by such models is, at the very least, a resource we use in both thinking and communication. Importantly though, this project can be run without having to decide on the semantic status of such representations.

It is highly controversial whether prototype-theoretical representations should feature in our formal representations of lexical meanings and concepts. At least some reason for thinking they should not, is that prototype structures face challenges when it comes to compositionality. To some extent, however, a failure of compositionality would be more serious for prototype accounts, given that, even if not representations of meanings *per se*, an adequate theory of communication would still require the information represented by such structures to be compositional.

At least for Boolean combinations, I have provided some reason to think that this requirement can be met by interpreting prototype structures in a Bayesian way. There is still a problem, however, of whether it is adequate to allow non-Boolean combinations of prototype structures to combine in a context-sensitive way, and what the mechanisms for such combinations would be on a Bayesian model.[9]

Finally, the model provided here should be tested against a wide range of Boolean $N \cdot N$ combinations to determine what kind of coverage this Bayesian approach is able to attain with respect to empirically sourced typicality data.

References

Austin, J. L., 1950/1979: Truth. In J. O. Urmson and G. J. Warnock (eds.), *Philosophical Papers, Third Edition*, Oxford: Oxford University Press. 117–133. Originally in: *Symposium: Truth, Proceedings of the Aristotelian Society*, Vol. 24 (1950).

[9]Howeverm an analysis of $A \cdot N$ combinations in a Bayesian situation-theoretic model has been detailed in (Author citation, year).

Austin, J. L., 1970: The Meaning of a Word. In J. O. Urmson and G. J. Warnock (eds.), *Philosophical Papers, Second Edition*, Oxford: Oxford University Press. 54–75.

Barsalou, L. W., 1992: Frames, concepts, and conceptual fields. In *Frames, Fields, and Contrasts: New essays in semantic and lexical organization*, Lawrence Erlbaum Associates. 21–74.

Borg, E., 2004: *Minimal Semantics*. Oxford University Press.

Cappelen, H. and Lepore, E., 2005: *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Oxford: Blackwell.

Connolly, A. C., Fodor, J. A., Gleitman, L. R. and Gleitman, H., 2007: Why Stereotypes Don't Even Make Good Defaults. *Cognition*, 103(1), 1–22.

Cooper, R., 2010: Frames in Formal Semantics. In H. Loftsson, E. Rgnvaldsson and S. Helgadttir (eds.), *Advances in Natural Language Processing*, Springer Berlin Heidelberg, volume 6233 of *Lecture Notes in Computer Science*. 103–114.

Cooper, R., 2012: Type Theory and Semantics in Flux. In R. Kempson, T. Fernando and N. Asher (eds.), *Philosophy of Linguistics, Handbook of the Philosophy of Science*, Elsevier. 271–323.

Cooper, R., Dobnik, S., Lappin, S. and Larsson, S., 2013: Probabilistic Type Theory and Natural Language Semantics. *ms, University of Gothenburg and King's College London.*

Cooper, R., Dobnik, S., Lappin, S. and Larsson, S., 2014: A Probabilistic Rich Type Theory for Semantic Interpretation. *Proceedings of the EACL 2014 Workshop on Type Theory and Natural Language Semantics.*

Dobnik, S., Cooper, R. and Larsson, S., 2013: Modelling Language, Action, and Perception in Type Theory with Records. In D. Duchier and Y. Parmentier (eds.), *Proceedings of the 7th International Workshop on Constraint Solving and Language Processing (CSLP 2012)*, Springer-Verlag. 70–91.

Fillmore, C. J., 1975: An Alternative to Checklist Theories of Meaning. *Proceedings of the First Annual Meeting of the Berkeley Linguistics Society*, 1, 123–131.

Fillmore, C. J., 1976: Frame semantics and the nature of language. *Annals of the New York Academy of Sciences*, 280(1), 20–32.

Fillmore, C. J., 1982: Frame semantics. In *Linguistics in the Morning Calm*, Hanshin. 111–137.

Fodor, J. A., 1998: *Concepts: Where Cognitive Science Went Wrong*. 438. Oxford University Press.

Fodor, J. A. and LePore, E., 1991: Why Meaning (Probably) Isn't Conceptual Role. *Mind and Language*, 6(4), 328–43.

Fodor, J. A. and LePore, E., 1996: The Red Herring and the Pet Fish: Why Concepts Still Can't Be Prototypes. *Cognition*, 58(2), 253–70.

Hampton, J. A., 1991: The combination of prototype concepts. In P. Schwanenflugel (ed.), *The psychology of word meanings*, Lawrence Erlbaum. 91–116.

Hampton, J. A., 1995: Testing the prototypical theory of concepts. *Journal of memory and language*, 34, 686–708.

Hampton, J. A., 2007: Typicality, Graded Membership, and Vagueness. *Cognitive Science*, 31(3), 355–384.

Hampton, J. A. and Jönsson, M. L., 2012: Typicality and Compositionality: The logic of combining vague concepts. In Hinzen, Machery and Werning (eds.), *The Oxford Handbook of Compositionality*, Oxford University Press. 386–401.

Kamp, H. and Partee, B., 1995: Prototype Theory and Compositionality. *Cognition*, 57(2), 129–191.

Kennedy, C., 1997: *Projecting the adjective: the syntax and semantics of gradability and comparison. PhD Thesis*. University of California, Santa Cruz.

Kennedy, C., 2007: Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives. *Linguistics and Philosophy*, 30(1), 1–45.

Lewis, D., 1970: General Semantics. *Synthese*, 22(1-2), 18–67.

Osherson, D. and Smith, E. E., 1982: Gradedness and Conceptual Combination. *Cognition*, 12, 299–318.

Osherson, D. and Smith, E. E., 1997: On Typicality and Vagueness. *Cognition*, 64(2), 189–206.

Prinz, J. J., 2002: *Furnishing the mind: concepts and their perceptual basis*. Mit Press.

Prinz, J. J., 2012: Regaining Composure: A Defence Of Prototype Compositionality. In Hinzen, Machery and Werning (eds.), *The Oxford Handbook of Compositionality*, Oxford University Press. 437–453.

Putnam, H., 1975: The Meaning of 'Meaning'. *Minnesota Studies in the Philosophy of Science*, 7, 131–193.

Rayo, A., 2011: A Plea for Semantic Localism. *Noûs*, 47(4), 647–679.

Recanati, F., 2004: *Literal Meaning*. Cambridge University Press.

Rey, G., 1983: Concepts and Stereotypes. *Cognition*, 15, 237–262.

Robbins, P., 2002: How to Blunt the Sword of Compositionality. *Noûs*, 36(2), 313–334.

Rosch, E., 1978: Principles of Categorization. In *Cognition and Categorization*, Lawrence Elbaum Associates. 27–48.

Rosch, E. and Mervis, C. B., 1975: Family Resemblances: Studies in the Internal Structure of Categories. *Cognitive Psychology*, 7(4), 573–605.

Smith, E. E., Osherson, D., Rips, L. J. and Keane, M., 1988: Combining Prototypes: A selective modification model. *Cognitive Science*, 12, 485–527.

Thomason, R. H. (ed.), 1974: *Formal Philosophy, Selected Papers of Richard Montague.* Yale University Press.

Travis, C., 1978: Meaning Versus Truth. *Dialogue*, 17(03), 401–430.

Wisniewski, E. J., 1997: When concepts combine. *Psychonomic Bulletin & Review*, 4(2), 167–183.

Wittgenstein, L., 1953/2009: *Philosophical Investigations* Revised 4$^{\text{th}}$ Edition, Hacker and Schulte (eds.). Oxford: Blackwell.