Polysemy in Conceptual Combination: Testing the Constraint Theory of Combination

Fintan Costello & Mark T. Keane

Dept. of Computer Science,

University of Dublin, Trinity College,

Dublin 2, IRELAND

`fintan.costello, mark.keane@cs.tcd.ie`

Abstract

Most novel noun-noun combinations are polysemous in that they tend to suggest several possible meanings. A *finger cup* can be a cup in which fingers are washed, a cup shaped like a finger, a narrow cup and so on. In this paper, we present a new theory of concept combination, the constraint theory, that accounts for the polysemy of noun-noun combinations. Constraint theory, which uses three constraints (diagnosticity, plausibility and informativeness) acting over a unitary mechanism that generates candidate interpretations, makes certain predictions about the polysemy of different combinations. In particular, it predicts that combinations involving artifact terms should be more polysemous than those involving natural kinds because the former have functional models that promote multiple interpretations. In a single experiment, this prediction is confirmed along with other predictions about the types of interpretation that tend to be produced.

Introduction

Research on concept combination lies at the heart of our understanding of natural language and concepts (see Gleitman & Gleitman, 1970; Murphy & Medin, 1985; Osherson & Smith, 1981). Most previous research has concentrated on how people arrive at a single interpretation for a *modifier-head* combination like *car repair* or *plastic dog* (see Murphy, 1988, 1990; Wisniewski, 1995; but see Kay & Zimmer, 1976). This research has tended to underplay the creativity of conceptual combination, and how inventive the products of combination can be. For concept combination can be viewed, like analogy, as a method for creating new knowledge from previous experience (see e.g., Keane, in press; Keane, Ledgeway & Duff, 1994). Taking this perspective, we look at the *polysemy* of conceptual combination, at the multiple interpretations an individual can produce to a single novel noun-noun combination [1].

---

[1] We assume that ambiguous words in combination will affect polysemy (e.g., bat ball); so, our concern is with how the conceptual representation of two normally unambiguous words affects polysemy (in Wisniewski & Gentner's, 1991, terms we examine the *relational ambiguity* not *lexical ambiguity* of combinations).

So, typically, we present subjects with novel noun-noun combinations and note the frequency and variety of interpretations they produce to each combination. Table 1 shows just some of the interpretations produced to the novel combination *bed pencil*.

Table 1: Interpretations for *bed pencil*

---

(1)	a pencil that you put beside your bed for writing some messages
(2)	a pencil used to draw patterns on bed-clothes
(3)	a bed made out of pencils
(4)	a pencil shaped like a bed
(5)	a thin bed
(6)	a big, flat pencil that is a bed for a doll
(7)	a pencil case

The present research is motivated by several aims. First, to examine several little-investigated factors that affect polysemy. Second, to determine whether the results from single-interpretation paradigms extend to this multiple-interpretation one. Third, to learn about the structure of our conceptual systems. Fourth, to determine the baseline ambiguity in combinations that is resolved by context (see e.g., Murphy, 1990). In the following sections, we review some of the empirical and theoretical work in this area before presenting a new theory of combination and a test of this theory.

Some Empirical Findings in Concept Combination

Single-interpretation research has established several regularities about the content of interpretations (see e.g., Hampton, 1987; Murphy, 1988; Smith, Osherson, Rips & Keane, 1988; Wisniewski & Gentner, 1991). Three main categories of interpretation have been studied (typical percentages based on single-interpretation studies by Wisniewski and colleagues are indicated):

- *Relational interpretations* (30%-50%) establish some relationship between the modifier concept and the head concept (e.g., 1, 2 and 3 in Table 1)
- *Property-mapping interpretations* (30%-50%) assert a property of one concept as true of the other concept (e.g., 4 and 5 in Table 1)
- *Hybrid interpretations* (0-10%) which are *both* the modifier-concept and the head-concept (e.g., "a colorful fish that is kept as a pet" for *pet fish*.; closest example in Table 1 is 6)

To these we add a fourth category that has not been closely examined before:

- *Direct-reference interpretations* (?%) where the interpretation is a known concept (e.g., 7 in Table 1)

The second main empirical regularity is that the *referent* of the interpretation (what it is about) is typically the 2nd noun or *linguistic-head* of the combination following the pragmatics of English noun-noun phrases (Gleitman & Gleitman, 1970; e.g., see 1, 2, 4, 6 in Table 1). However, in some cases, *reversals* occur in which the 1st noun or *linguistic-modifier* is the referent (e.g., 3 and 5 in Table 1). Wisniewski and Gentner (1991) found such reversals in substance-count phases (e.g., a *plastic dog* is a piece of plastic shaped like a dog). Finally, the referent can even be a concept other than those mentioned in the combination, which are termed *exocentric* in linguistics (Cannon, 1987; e.g., 7 in Table 1). It remains to be seen whether these empirical regularities hold when we consider the multiple interpretations produced by a single individual to a given combination.

Some Theories of Combination

Two main theories have concerned themselves with the above empirical regularities; concept-specialisation theory (e.g., Cohen & Murphy, 1984; Murphy, 1988) and alignment theory (e.g., Wisniewski & Markman, 1993; Markman & Wisniewski, in press).

Concept-specialisation theory maintains that the combination process specialises a slot of the head concept with the modifier concept. So, " a pencil used in bed" specialises the function-slot of bed with pencil (see Table 1). This theory expects relational interpretations with linguistic-head referents. However, it does not easily predict the other interpretation-types or reversals (i.e., it is not clear whether they could emerge from the theory's ill-specified elaboration process).

Alignment theory maintains that combination involves an alignment process, in which the concepts' slots are aligned to find commonalties (matching slot-value pairs) and alignable differences (matching slots with different values). These alignable differences are then used to identify properties to map (e.g., in *elephant bird* the size-slot matches but the values differ, allowing BIG to be mapped onto bird). On its own, this theory expects property-mapping and hybrid interpretations to be produced and

seems to allow reversals to occur (see Wisniewski & Gentner, 1991). However, it is silent on direct-reference interpretations and requires a specialisation mechanism to deal with relational interpretations.

Both of these theories concentrate on explaining how a single, acceptable meaning for a compound is produced. In contrast, a theory of polysemy must account for the variety of meanings that are produced to a given compound and how these interpretations might vary in their relative acceptability. We propose a new *constraint theory* of concept combination that explicitly deals with polysemy using a unitary mechanism to explain all the interpretation-types, which makes specific predictions about the nature of reversals.

Constraint Theory of Concept Combination

Constraint theory has three components: (i) a generative mechanism that produces a set of candidate interpretations (ii) three constraints -- diagnosticity, plausibility and informativeness -- that decide between these interpretations and (iii) a set of knowledge assumptions about the types of concepts used in the combination.

The Generative Mechanism

The generative mechanism forms all subsets of predicates of the constituent concepts, along with predicates from knowledge related to these concepts (see Costello, 1996). This unitary mechanism generates a wide range of candidate interpretations, which can have a variety of referents. Note that even though this mechanism can produce relational and property-mapping interpretations, it does so without using specialisation or alignment mechanisms.

The Three Constraints

At the heart of the theory, three constraints operate over the candidate interpretations produced. These constraints can reject some interpretations and/or promote some as being more acceptable than others; hence, dealing directly with the issue of relative acceptability.

The *diagnosticity constraint* requires interpretations to contain diagnostic predicates of both concepts. In our terms, a diagnostic predicate is one that strongly distinguishes the concept from other known, related concepts (similar to Rosch's, 1978, *cue validity* idea). Diagnosticity makes the second interpretation more acceptable than the first:

- a cactus fish is a green fish
- a cactus fish is a prickly fish

because PRICKLY is more diagnostic of cactus than GREEN. The diagnosticity constraint also requires that *the referent of the interpretation possess diagnostic predicates of the linguistic-head* (i.e., the 2nd noun). This does not disallow reversals but requires these linguistic-modifier referents to have the diagnostic predicates of the linguistic-

head. Usually, this will mean that diagnostic properties of the linguistic-head will be asserted of the linguistic-modifier. For example, Wisniewski & Gentner (1991) report a reversal interpretation for *chair ladder* of "a chair that is by necessity used as a ladder" where the diagnostic function of *ladder* is asserted of the *chair*. Thus, the theory makes the novel prediction that reversals should tend to be property-mapping interpretations (a prediction we test later).

The *plausibility constraint* requires that interpretations describe an object (or collection of objects) which could plausibly exist. Plausibility makes the second interpretation more acceptable than the first:

• an angel pig is a pig with wings on its tail
• an angel pig is a pig with wings on its torso

because prior experience dictates that wings are typically attached to the centre of gravity of an object (excepting Hermes). Plausibility is computed from the degree to which the semantic elements of an interpretation have occurred together in known concepts. Plausibility predicts direct-reference interpretations because, being known concepts, they have the highest possible plausibility score.

The *informativeness constraint* requires interpretations to communicate something new (e.g., Grice, 1975). Informativeness excludes feasible interpretations that do not communicate anything new relative to either constituent concept. So, it would predict that people should not produce the following:

• A bed pencil is a pencil made of wood

because MADE-OF-WOOD does not convey any new information about either *beds* or *pencils*. It also accounts for Downing's (1977) finding that people find combinations with *redundant modifiers* unacceptable; for example, pig pork. Informativeness plays an important role in deciding between interpretations which tie on diagnosticity and plausibility.

Types of Knowledge

Constraint theory also has assumptions about the types of concepts used in combinations. As in the above theories, we assume that concepts have complex representational, predicate structure (consisting of attributes, objects and relations). In this paper, we examine an assumption about the basic ontological distinction between artifacts and natural kinds (see e.g., Keil, 1986).

Artifact concepts (*guns, buildings,* and *cups*) differ in several ways from natural kinds (*birds*, *trees*, and *snails*); predominantly, in that artifacts are functionally-related to other concepts. These functional models should facilitate polysemy because they typically have several possible roles (e.g., agent, object, recipient, instrument). For example, an *elephant gun* can be a gun used by an elephant to shoot things (agent role) or a gun used to shoot elephants (object role). Natural kinds do not normally have associated functional models and therefore should support less polysemy. Furthermore, the relational roles in artifacts tend to have a wider scope those of natural kinds, in the sense that a wide range of objects can be used to fill them (see Wisniewski & Gentner, 1991). A gun can be used to *shoot* almost anything but a snail can only *eat* vegetative matter.

Both of these factors, along with the diagnosticity constraint's requirement that the referent of an interpretation possess diagnostic properties of the linguistic-head, suggest that combinations with an artifact head should manifest significantly more polysemy than those with a natural-kind head.

Concept-Type & Polysemy

This experiment tested three main predictions of constraint theory. First, the prediction that all four types of interpretations should occur (notably the direct-reference predicted by plausibility). Second, that when reversals occur they should tend to be property-mappings (based on diagnosticity). Third, that concept-type should affect the polysemy of combinations; specifically, that artifact-head combinations should generate more interpretations than natural-kind head ones (based on more functional interpretations).

Subjects were given four sets of six combinations from the following categories: artifact~artifact (e.g., "pencil bed"), natural-kind~artifact (e.g., "river chair"), artifact~natural-kind (e.g., "chair river") or natural-kind~natural-kind combinations (e.g., "oak dog").

Method

Subjects & Design. Twenty undergraduates at Trinity College Dublin took part voluntarily in the experiment (one was excluded prior to data analysis). The design was a 2 (head-category) x 2 (modifier-category) one, both being within-subjects variables. Combinations were counterbalanced for the order of the specific words used (e.g., each subject got *bed pencil* and *pencil bed*). This variable which is not reported here does not change the effects found (see Costello, 1996, for details).

Materials. Twenty-four basic-level words (12 artifact words and 12 natural-kind words) were used to construct the 24 combinations (see Appendix). Frequency of occurrence of the words in each set determined using the Oxford Psycholinguistic Database (Quinlan, 1992). A t-test comparing artifact and natural-kind word frequency showed no reliable difference between the categories ($t(22) = .84$, $p > .4$).

Procedure. The instructions asked subjects to say "what the phrase could plausibly mean and if you can think of more than one possible meaning for a phrase then report them in the order in which they occur to you", for the novel combinations they were shown. The combinations were presented individually on cards and subjects' responses were audio recorded. Subjects were tested individually.

Scoring. Subjects' interpretations were classified by both experimenters into the four interpretation-types[2]: relational, property-mapping, hybrid and direct reference. They were also classified by referent: head, reversal or other. Finally, the frequency of functional interpretations were noted. Differences were resolved by discussion (agreement was 97%).

Results & Discussion

The results confirmed the main predictions of the constraint theory in finding that head-artifact combinations were more polysemous than head-natural kinds ones, based on increased numbers of functional interpretations being produced in the former.

Types of Interpretation. Each of the four types of interpretation were found among the 1019 produced: relational 46% (468), property-mapping 33% (337), hybrid 0.3% (3), and direct-reference 15% (157), with 6% (57) in the other category. These results are in line with previous findings on the first three categories and confirm that direct-referent interpretations are a commonly produced type. The findings support the flexible generative mechanism used in the constraint theory and its proposed constraints.

Polysemy. On average each subject produced about two interpretations per combination ($M = 2.24$), although the frequency ranged from 1 to 6. The 2 head-category (artifact or natural-kind) x 2 modifier-category (artifact or natural-kind) repeated-measures ANOVA revealed a main effect of head-category; with artifact-head combinations producing reliably more interpretations ($M = 2.36$) than natural-kind-head ones ($M = 2.12$,; $F(1,17) = 7.89$, $p =.01$, $MSe = .78$). No other reliable effects or interactions were found. An item analysis revealed similar effects of head-category ($F(1,22) =5.71$, $p < .05$, $MSe = 1.12$).

An analysis of the percentage of functional interpretations produced to a given phrase showed that there was a higher percentage of functional interpretations produced to artifact-head ($M = .40$) than to natural-kind-head combinations ($M =.12$; $F(1,17) = 72.96$, $MSe = .122$, $p < .0001$). A similar, but lesser, difference was found between artifact-modifier ($M= .33$) and natural-kind-modifier combinations ($M =.19$; $F(1,17)= 18.23$, MSe = .115, $p < .0005$).

Both of these results support the core assumptions of the constraint theory. In particular, they show that artifacts are much more polysemous than natural kinds, especially when they are in the head position. Furthermore, it shows that this extra polysemy is based on their associated functional models.

Referents & Reversals. As expected, most of the referents of interpretations were to the linguistic-head of the combination (68%), but reversals using the linguistic-modifier as the referent were relatively common (9%), with the remainder having some other concept as the referent (23%). This result confirms that such reversals are not limited to combinations of a substance-object form but occur more generally, albeit not as frequently (Wisniewski & Gentner, 1991, found rates as high as 38%).

The diagnosticity constraint predicts that when reversals occur they should tend to be property mappings (of the head's diagnostic predicates) rather than another interpretation-type. This prediction is confirmed by an analysis of interpretation types in the different referent categories. In the linguistic-head category, the interpretations produced tend to be relational (51%) rather then property-mappings (36%; a similar pattern is found in the Other category). In contrast, in the reversals category interpretations tend to be property mappings (54%) rather than relational ones (31%). All these differences are statistically reliable.

Significance of Findings. With any new paradigm the issue of the relevance of the findings always arises. Does the process by which people produce multiple interpretations have anything to do with people's normal comprehension of novel compounds? In particular, do contextual influences normally banish the possibility of polysemy in novel combinations? We would like to argue that the combination process used to produce multiple interpretations is the same as that which produces a single interpretation, with contextual influences acting as an extra constraint in this process.

In support of this position it may be noted that the types of interpretations produced in the present study and the relative frequencies of these interpretation-types closely parallel those found in traditional, single-interpretation studies. This is strong evidence for the proposition that the same process is being tapped in this and previous studies.

Granted, people may not normally entertain all the possible meanings reported by our subjects when a novel combination is presented in context (Murphy, 1990). For example, a *pencil bed* could be a narrow bed, a container for pencils, a bed that is pencil-shaped and so on. If you are told that "the pencil bed is in the bedroom upstairs" you are likely to assume that a narrow bed is upstairs, whereas if you are told that "the pencil bed is in the middle of the exam hall" you are more likely to think that it is some receptacle for pencils. However, not all contexts will necessarily disambiguate a novel combination; if you are told "he moved the pencil bed last week" either of the above two meanings could still hold (Mulligan, 1997). So, context will not necessarily banish all ambiguity.

Indeed, we believe that polysemy should be studied to tell us more about contextual influences in normal comprehension. At present we have little knowledge about the baseline ambiguity of novel combinations, and thus do not know what problems of ambiguity are resolved by context (see Mulligan, 1997).

---

[2] Operational definitions of the first three of these categories were similar to Wisniewski's (1996).

General Discussion

The present paper reports a new theory of conceptual combination that aims to explain specific aspects of the multiple interpretations produced by people to novel noun-noun combinations. Using this theory, we have shown that different types of concepts differentially affect polysemy; that combinations with artifact heads are more polysemous than those with natural-kind heads (similar results have been found for the superordinate/basic-level distinction). Furthermore, we have confirmed the generality of these findings in other experiments using several hundred different combinations (see Costello, 1996). We have also simulated these results in a model of the theory called C^3 (the Constraint Model of Concept Combination).

In constraint theory, we have tried to deal with the creative contents of interpretations produced by people to a single combination. At present, we have not concerned ourselves with more traditional concerns like typicality ratings (although ultimately we plan to extend the theory in this direction). In doing this we have tried to open up a new area of research on the factors that affect polysemy that appears to be very fruitful. In the course of this we have confirmed several ideas about the structure of the human conceptual system. Theoretically, we have been driven by the parsimony of using a single unitary mechanism and several simple explanatory constraints to account for these effects.

References

Cannon, G. (1987). *Historical change and English word-formation*. New York: Lang.

Cohen, B. & Murphy, G. L. (1984). Models of concepts. *Cognitive Science*, 8, 27-58.

Costello, F. (1996). *Noun-noun conceptual combination: the polysemy of compound phrases*. PhD Thesis University of Dublin, Dublin, Ireland.

Downing, P. (1977). On the creation and use of English compound nouns. *Language, 53*, 810-842.

Gleitman, L.R. & Gleitman, H. (1970). *Phrase and paraphrase*. New York: Norton.

Grice, H.P. (1975). Logic and conversation. In P. Cole and J.L. Morgan (Eds.), *Syntax and semantics (vol 3): Speech acts*. New York: Academic Press.

Hampton, J. A. (1987). Inheritance of attributes in natural concept conjunctions. *Memory and Cognition, 15*, 55-71.

Kay, P. & Zimmer, K. (1976). On the semantics of compounds and genitives in English. *Sixth California Linguistics Association Proceedings.* San Diego, CA: Campile Press.

Keane, M.T. (in press). What makes an analogy difficult ?: The effects of order and causal structure in analogical mapping. *Journal of Experimental Psychology: Language, Memory & Cognition.*

Keane, M.T., Ledgeway, T, & Duff, S. (1994). Constraints on analogical mapping: A comparison of three models. *Cognitive Science*, *18*, 287 - 334.

Keil, F. (1986). The acquisition of natural kind and artifact terms. In W. Demopoulos & A. Marras (Eds.), *Language learning and concept acquisition.* Norwood, NJ: Ablex.

Rosch, E. (1978). Principles of categorisation. In E. Rosch & B. B. Lloyd (Eds.) *Cognition and categorisation.* Hillsdale, NJ: Erlbaum.

Markman, A.B., & Wisniewski, E. J. (in press). Same and different: The differentiation of basic-level categories. *Journal of Experimental Psychology: Language, Memory & Cognition.*

Mulligan, M. (1997). *Context & concept combination.* University of Dublin, Trinity College, Technical Report TCD-CS-97-04.

Murphy, G. L. (1988). Comprehending complex concepts. *Cognitive Science*, *12*, 529-562.

Murphy, G. L. (1990). Noun phrase interpretation and conceptual combination. *Journal of Memory and Language, 29*, 259-288.

Murphy, G. L. & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, *92*, 289-316.

Osherson, D. N. & Smith, E. E. (1981). On the adequacy of prototype theory as a theory of concepts. *Cognition, 9*, 35-58.

Smith, E. E., Osherson, D. N., Rips, L. J. & Keane, M. (1988). Combining prototypes: A selective modification model. *Cognitive Science*, *12*, 485-527.

Quinlan, P. T. (1992). *The Oxford psycholinguistic database.* Oxford: Oxford University Press.

Wisniewski, E. J. & Gentner, D. (1991). On the combinatorial semantics of noun pairs: Minor and major adjustments to meaning. In G. B. Simpson (Ed.) *Understanding word and sentence.* Amsterdam: North Holland.

Wisniewski, E. J. & Markman, A. B. (1993). The role of structural alignment in conceptual combination. *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society.* Boulder, CO.

Wisniewski, E. J. (1996). Construal and similarity in conceptual combination. *Journal of Memory and Language, 35*, 434-453.

Appendix:
Combinations Used in the Experiment

artifact-artifact	Bed Pencil Bus Shirt Train Hat	Pencil Bed Shirt Bus Hat Train
artifact-natural~kind	Chair River Gun Horse Shoe Herring	River Chair Horse Gun Herring Shoe
natural-kind~artifact	Apple Knife Potato Ball Rose Hammer	Knife Apple Ball Potato Hammer Rose
natural-kind~natural-kind	Boulder Grass Eagle Tulip Oak Dog	Grass Boulder Tulip Eagle Dog Oak