

# The Iteration of Concept Combination in Sense Generation

**Richard Cooper**

Department of Psychology  
Birkbeck College, London  
Malet Street, London WC1E 7HX  
R.Cooper@psyc.bbk.ac.uk

**Bradley Franks**

Department of Psychology  
London School of Economics  
Houghton Street, London WC2A 2AE  
B.Franks@lse.ac.uk

## Abstract

We report work in progress on the computational modelling of a theory of concepts and concept combination. The *sense generation* approach to concepts provides a perspicuous way of treating a range of recalcitrant concept combinations: privative combinations (e.g., *fake gun*, *stone lion*, *apparent friend*). We argue that a proper treatment of concept combination must respect important syntactic constraints on the combination process, the simplest being the priority of syntactic modifier over the head in case of conflicts. We present a model of privative concept combinations based on the sense generation approach. The model was developed using COGENT, an object-oriented modelling environment designed to simplify and clarify the implementation process by minimising the 'distance' between the box/arrow 'language' of psychological theorising and the theory's implementation. In addition to simple privatives (i.e., ones with a single modifier, like *fake gun*) the model also handles iterated, or complex, privative combinations (i.e., ones with more than one modifier, like *fake stone lion*), and reflects their associated modification ambiguities. We suggest that the success of this model reflects both the utility of COGENT as a modelling framework and the adequacy of sense generation as a theory of concept combination.

## Introduction

Concepts are usually taken to have three basic functions in mental life (see, e.g., Franks, 1992; Rips, 1995): a *representation* function (the representational contents over which thought and inference about the world takes place), a *classification* (or referring) function (the contents employed in determining whether objects fall under the denotation of a term), and a *linguistic* function (the contents accessed in understanding language and concatenated according to linguistically appropriate rules to comprise a mental representation of the meaning of a sentence or utterance). The first two functions have been the primary concern of investigations of concepts in cognitive science.

It is widely accepted that insights about the representation and classification functions of concepts can be obtained from an understanding of the way in which they combine to form complex concepts (e.g., Medin & Shoben, 1988; Smith, Osherson, Rips & Keane, 1988). We suggest that constraints on the process of combination itself should be forthcoming from an understanding of the linguistic function. Hence, accounts of concepts in general can be constrained by ascertaining the extent to which they respect critical factors concerning the linguistic function. In order to begin to make progress on this question, the work reported here focuses on the effects of

aspects of the linguistic function on the representational function of concepts (since, arguably, the classification function is parasitic on the representational function).

We present a computational model of concept combination from within the sense generation framework developed elsewhere (e.g., Franks, 1995), employing a computational framework developed in order to clarify the relationship between the specification of theoretical commitments and their implementation (Cooper, Fox, Farrington & Shallice, in press). This is an essentially symbolic framework for modelling cognitive phenomena: since our focus was on modelling the effect of syntactic relations on concept combination (rather than modelling those syntactic relations *per se*), the necessary perspicuity and clarity of the sequential ordering in the implementation of the syntax is more directly lent by a symbolic implementation than a connectionist one (such as Miikkulainen, 1993). We present a basic model of concept combinations first (a head noun combining with a single modifier), and then show how this can be extended to handle a simple form of syntactic influence on combination — the iteration of modifiers, with attendant scoping or modification ambiguities that arise for their representation. The model presented is an aspect of work in progress. This work has a goal of locating a model of concept combination within a perspective gained from syntactic constraints on combination and from wider considerations of cognitive modelling.

## Sense Generation and Concept Combination

Sense generation is an approach to concepts and concept combination that attempts to respect psychological evidence about classification and representation in the context of pragmatic factors concerning communication. It postulates 'quasi-classical' *lexical concept* representations which express the default content for a concept, comprising attribute-value structures where each attribute takes only one value, and where those values can be overridden by contextual dictates (Franks & Braisby, 1990; Franks, 1995). These representations are quasi-classical in that, although they act as if they were necessary and sufficient conditions for category membership *within* a single context, *across* different contexts their contents are defeasible and so are not classical definitions. A critical distinction between types of attributes is made on the basis of their relationship to category membership: 'central' attributes are ones that reflect deep, theoretical assumptions about the factors that are presumed essential for category membership (see Medin & Ortony, 1989; Braisby, Franks & Hampton, 1994, in press); by contrast, 'diagnostic' attributes are those aspects of

the surface appearance of objects that may be used for rough-and-ready identification, but are not infallible guides. Sense generation postulates a generative process that takes bottom-up input from lexical concepts associated with the constituents of a phrase, and outputs a *sense* for the phrase that provides a closer fit to the pragmatic context than would the default content of lexical concepts.

It can be argued (Franks, 1995) that a class of combinations known as 'privatives' constitute a test-case for theories of concept combination, in that they exhibit particularly strong, yet predictable, forms of context-sensitivity. Privative adjectives are analysed by Kamp (1975) as ones for which the following inference is a logical truth: *If  $x$  is a privative adjective-noun then  $x$  is not a noun*: for example, if  $x$  is a *fake gun*, then  $x$  is not a *gun*.

Franks (1995) argues that this analysis should be extended in three ways. First, the characteristic inference should be weakened, to allow for the cognitively plausible classification, if  $x$  is a *fake gun*, then  $x$  is a *gun* in some sense — only with respect to appearance. Second, such an inference is characteristic not only of the effect of a particular set of adjectives on all nouns that they modify, but also of the interaction of the contents of head nouns with modifiers that are not themselves intrinsically privative. Third, 'proper' privatives, in which privative behaviour results from an intrinsic property of the adjective type, in fact come in two kinds — 'negating' privatives (e.g., *fake gun* — where, intuitively, the modifier negates the central attributes of the head), and 'equivocating' privatives (e.g., *apparent friend* — where the modifier casts doubt on the head's central attributes, but does not completely negate them). In both cases, the diagnostic attributes of the head noun concept are not denied in any way — thus preserving the characteristic classification inference based on appearances only, noted above. These types of privative have analogues in which privative behaviour results not from the particular semantic type of the modifier, but from the interaction of the contents of the concepts of head and modifier. Such 'functional' privatives include combinations like *stone lion* or *chocolate teapot* (negating privative analogues of *fake gun*), and *wooden skillet* or *blue orange* (equivocating privative analogues of *apparent friend*). It is clear that, for these cases, there is no intrinsic property of either the head or modifier that produces the privative behaviour (e.g., when *stone* modifies *bridge*, and when *lion* is modified by *brown*, the resulting behaviour is not privative). Hence, privatives embody a particularly strong form of context-sensitivity of concept representations. Producing the sense for a privative combination denies attributes of the head noun, requiring a process that is not rigidly (i.e., monotonically) compositional (for example, as in feature-addition; see Hampton, 1987). This does not preclude a compositional account, however, since the sense produced for a combination is still predictable from the concepts of the parts and their mode of combination.

Despite their constituting a test case for theories of concepts and concept combination, it is not clear that privatives can be handled in, for example, prototype theory (Hampton, 1992), schema theory (Murphy, 1988), or theories that assume that combination operates by a process of either property mapping or slot filling (cf., Wisniewski & Gentner, 1991). Moreover, the requisite marked conflict between attributes of the head

and modifier may also be difficult to express in connectionist terms. By contrast, a unified, compositional account of privatives within the sense generation framework is presented in Franks (1995). The account employs unification-based operations to capture the various aspects of concept combination (e.g., priority unification (Kaplan, 1987), in which the sense for the combination inherits all of the attribute-value pairs of the constituent lexical concept, with the exception that where the two concepts conflict on one of the values of a common attributes, the value of the priority constituent — the modifier — is inherited). The critical difference between negating and equivocating privatives is expressed in terms of different metonymic type coercion operators (cf., Klein & Sag, 1985; Pustejovsky, 1991): for negators, the operator takes as input the head noun's lexical concept (comprising both central and diagnostic attributes), and outputs a coerced representation comprising the head's diagnostic attributes and the negation of its central attributes; for equivocators, the operator outputs the diagnostic attributes of the head and neither the central attributes nor their negation, but rather their being only possibly true of the object being referred to. This corresponds to treating the attribute-value structures as potentially having three truth-values — a value of an attribute may be true of a type of object, false of it, or neither. For functional privatives, the conflict between central attributes for *stone* and *lion* results in those of the latter being negated, whilst the conflict between diagnostic attributes for *wood* and *skillet* results in the central attributes of the latter not being negated but being only possibly true of the type of object described by the phrase.

This treatment captures defining semantic intuitions about objects described by privatives, and hence their characteristic classification inferences. For example, a stone lion possesses central attributes of stone objects, but only diagnostic attributes of real lions. Similarly, fake guns do not possess the central attributes of guns, but only their diagnostic attributes. By contrast, an apparent friend performs diagnostic behaviours of a friend, but this is consistent with only *possibly* possessing central motivational attributes of friendship; similarly, a wooden skillet looks like a real skillet but may or may not be able to support the central function of a skillet. In all cases, the initial, bottom-up combination stage leaves the elaboration and specification of the details of the combination unspecified, producing schematic senses which are consistent with a range of further possible specifications that depend on informational and pragmatic context (e.g., resulting in the addition of the information that the stone lion is a statue, and so has central attributes of a statue, or that an apparent friend really is or is not a friend, with appropriate central attributes).

The model discussed here seeks to test this account in two ways. First of all, to ascertain whether the account of simple privative concept combinations (i.e., ones in which there is a single modifier for the head) is coherent by modelling it. Secondly, to ascertain whether the particular combination operations are adequate for handling iterated or complex concept combinations (in which there is more than one modifier). This also provides a way of beginning to incorporate syntactic constraints into concept combination.

Noun phrases often incorporate multiple modifiers of the head noun. The multiplication of modifiers raises the possibility of ambiguity in the scope of the modification, and this

problem is rendered especially complex when the multiplied modifiers are privatives, and hence when a privative combination is itself modified by a privative. For example, *fake stone lion* could have two distinct readings, one in which the first modifier has 'wide' scope over the second modifier and the head (i.e., *fake (stone lion)*: a stone lion that is a fake), and one in which the first modifier has 'narrow' scope over the second modifier only, and they both modify the head (i.e., *(fake stone) lion*: a lion that is made of fake stone).

In essence, in order to arrive at an appropriate representation for such complex combinations, the head-modifier relationships — that is, the scope of the first modifier — must be disambiguated. The disambiguation of such relationships is taken to be provided by a syntax parser that provides input to a conceptual interpreter. This does not imply, of course, that no semantic or thematic lexical information is involved in the process of syntactic disambiguation, merely that detailed conceptual information is not (see, e.g., Trueswell, Tanenhaus & Kello, 1993). The question of interest for our purposes is then, given that the head-modifier relationships have been disambiguated, does the sense generation account of simple privative combinations generalise to complex privatives?

### The COGENT Modelling Environment

Our model was constructed in COGENT (Cognitive Objects within a Graphical EnvironMent), a developing cognitive modelling tool. This system grew out of independent work aimed at clarifying the relationship between psychological theory and computational implementation within cognitive modelling. In particular, it seeks to make explicit the range of actual architectural commitments made by a psychological theory, as distinct from mere implementation details (cf. Cooper *et al.*, in press; Cooper, 1995).

The COGENT modelling environment provides a set of configurable cognitive "objects" (such as rule-based processes, and buffers with or without capacity limitations and content decay). Central to the environment is a graphical interface which allows the specification of executable models in the box/arrow style. Different shaped boxes correspond to different types of object, and a complete executable model can be developed by specifying appropriate properties and contents for all objects in the model. This minimises the distance between the models traditionally developed by psychologists and their implementations, simplifying the relation between the two. At present, only symbolic objects are provided, but anticipated future developments include extending the environment by incorporating connectionist and network objects. Extant uses include the implementation of production system models (Cooper, 1996), models of decision making (that of Fox, 1980), and models of prospective memory (Ellis, Shallice and Cooper, in submission). Details of COGENT availability and system requirements are available from the authors.

It is important to emphasise that COGENT is an environment, not an architecture. As such the intention is that the system should impose few (if any) constraints on the precise form of a model's implementation. In general, it is the particular theorist's responsibility to provide such constraints. This is not to say that architectures such as Soar (Newell, 1990) have no role in computational work on cognition (though see Cooper & Shallice (1995) for a discussion of some potential short-

comings of architecture-driven modelling), but to provide for theorists who do not subscribe to all of the assumptions embodied in any extant architecture.

In spite of the above intention, one constraint which is imposed by COGENT, and which impacts on the current domain, is the preclusion of recursive procedures (i.e., procedures which call themselves). Most current programming languages allow recursive procedures, and such procedures are of tremendous utility when processing tree-structured data such as natural language. However, true recursion requires a processing stack to maintain the trace of execution throughout recursive calls (in order to recover from the recursion once it bottoms out), and associated dynamic memory allocation for independent copies of local variables used within each recursive call.

While there is no prohibition on a COGENT process triggering itself, there is no processing stack within COGENT so there is no possibility that, on completion, a process could return control to the process that triggered it. (In fact, processing within COGENT is based on a parallel model in which all boxes are constantly potentially active: see Cooper (1995) for more details.) Furthermore, COGENT's assumed correspondence between cognitive objects (i.e., boxes) and functional cognitive structures means that, for example, a process requiring local variables will require an associated buffer in which to store those variables. A truly recursive process would require a separate copy of this buffer for each recursive call. Again, this dynamic memory allocation is not available within COGENT. Notice that the preclusion of recursion in COGENT arises not from any prior prejudice against recursive procedures, but from the directness of the mapping between functional units and COGENT objects, and from the underlying (parallel) processing model.

In light of the above, complex modified noun-phrases cannot be processed in COGENT by recursively processing their parts.

### Sense Generation in COGENT

The model consists of a set of interconnected boxes (see Figure 1) of two principal kinds: buffers and rule-based processes. A parsing process breaks input into a set of local trees (i.e., binary branching nodes which disambiguate head/modifier relationships) and adds its output to a temporary storage buffer (*Local Parse Trees*). We are not here concerned with the internal details of this parsing process. A second process, *Conceptual Access*, is triggered by elements in the store, processing them (i.e., accessing the conceptual content associated with their constituents) one at a time. These contents comprise the single-valued attribute value structures (with the major division between central and diagnostic attributes) noted earlier. The order of this processing is constrained by the requirement that the conceptual content of the constituents (i.e., head and modifier) is immediately available (either from the *Mental Lexicon* or from a short term store to which earlier processing may have written its results). As a consequence, processing is bottom-up.

Once the initial content for both the head and modifier has been obtained, the generation of the schematic sense for the phrase proceeds via two further processes. These are also controlled in a bottom-up manner, since they only depend upon the type of operator or upon whether conflicts occur

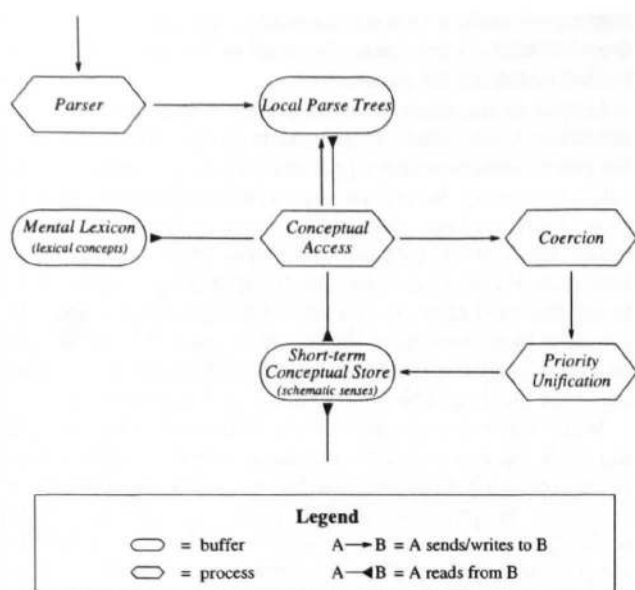


Figure 1: Sense Generation in COGENT

at diagnostic or central levels. The particular detailed contents of lexical concepts are not critical to the generation of a schematic sense. Firstly, the content of the head is coerced in one of two ways (see above): either the head's central attributes are negated (if the modifier is a negating privative adjective, or if the head and modifier conflict on central attributes), or the head's central attributes are undercut or cast into doubt (if the modifier is an equivocating privative adjective, or if the head and modifier conflict on diagnostic attributes). Following this, the coerced central and diagnostic attributes of the head are priority unified with the corresponding attributes of the modifier, with the modifier's attributes taking priority. This produces a schematic sense for the combination, which is temporarily stored in a conceptual store (cf., Potter, 1993). If this sense is associated with a constituent part of a complex noun phrase, then its arrival in the conceptual store will allow processing of its super-ordinate constituent to proceed: generating a sense for the complex noun phrase as a whole will then take inputs both from the *Mental Lexicon* (for lexical sub-constituents) and from the *Short-term Conceptual Store* (for non-lexical sub-constituents). The sense produced in this bottom-up manner is only schematic. Any post-combination specification or elaboration of this sense is viewed as involving an interaction between bottom-up and top-down influences and not modelled here.

The full specification of the model in COGENT consists of Figure 1 together with a specification of the configurable properties of each box in that figure. The theoretical differences between the *Mental Lexicon* and the *Short-term Conceptual Store*, for example, are reflected in different configurations of the corresponding boxes. In particular, the *Mental Lexicon* is modelled as a long term store with no decay and no capacity limitations, whereas the *Short-term Conceptual Store* is modelled as a temporary or working store with fixed decay and capacity limitations. Each process is fully specified in

terms of a set of condition/action rules (one rule per process for this model) and some declarative Prolog conditions. The conditions of the rules either match elements from various buffers or perform logical operations (such as priority unification, specified in Prolog) on their data. The rules' actions modify buffer contents or trigger further processes.

The three rules are presented in Figure 2. CH and DH represent the central and diagnostic attributes of the head (respectively). Similarly CM and DM represent the central and diagnostic attributes of the modifier. In Rule 1, the condition `conceptual_lookup` queries the *Mental Lexicon* and *Short-term Conceptual Store* to determine the content of the phrase's head. The condition `conflict_type` serves a similar purpose for the modifier, but also takes account of adjectives which serve as operators (and so have no independent conceptual content — e.g., proper privatives like *fake*). If these conditions are met, the rule fires, deleting the local parse tree from its buffer, and triggering *Coercion* with the message `coerce(Phrase, Type, CH, CM, DH, DM)`.

When Rule 2 is triggered, it coerces the type of the central attributes of the head according to that specified by the triggering process, producing CCH (i.e., coerced central head attributes). The process then triggers *Priority Unification* (Rule 3) with the message `unify_content(Phrase, CCH, CM, DH, DM)`.

The triggering of Rule 3 priority unifies the central attributes of the coerced head and modifier, producing the central attributes of the combination (C), and the diagnostic attributes of the head and modifier, producing the diagnostic attributes of the combination (D). The conceptual content of the combination is then added to the *Short-term Conceptual Store*, where it may contribute to the construction of the conceptual content of a larger constituent (via the condition `conceptual_lookup` called by *Conceptual Access* (Rule 1)).

To illustrate, consider first a simple privative, such as *stone lion*. There is no syntactic ambiguity, and just one local tree. Processing therefore requires just one cycle through the diagram in Figure 1. Once the local tree appears in the *Local Parse Trees* store, *Short-term Conceptual Access* is triggered, thus accessing the lexical concepts for both *stone* and *lion* from the *Mental Lexicon*: both lexical concepts comprise central and diagnostic attributes. *Conceptual Access* also determines the nature of the combination in terms of any conflict of attributes (i.e., negating or equivocating privative for conflicts on central or diagnostic attributes respectively, or affirmative combination for no conflict). In the case of *stone lion*, a conflict of central attributes will be detected, fulfilling the requirements for a negating type coercion. *Conceptual Access* passes this information to the *Coercion* process, which negates the central attributes of the head (*lion*). The output coerced representation then comprises the diagnostic attributes and the negated central attributes from the lexical concept for *lion*. The coerced representation of *lion* and the representation of the lexical concept for *stone* are then input to the *Priority Unification* process, which combines them (giving precedence to *stone*) to yield a representation of a type of object that possesses all of the central and diagnostic attributes of *stone*, and some of the coerced central and diagnostic attributes of *lion*, with the proviso that, where the values of attributes conflict, the values of *stone* take priority. The resulting representation

### Rule 1: Conceptual Access

IF: once local\_tree(Phrase, Modifier, Head) is in *Local Parse Trees*  
conceptual\_lookup(Head, CH, DH)  
conflict\_type(Modifier, Type, CH, DH, CM, DM)  
THEN: delete local\_tree(Phrase, Modifier, Head) from *Local Parse Trees*  
send coerce(Phrase, Type, CH, CM, DH, DM) to *Coercion*

### Rule 2: Coercion

TRIGGER: coerce(Phrase, Type, CH, CM, DH, DM)  
IF: type\_coerce(Type, CH, CCH)  
THEN: send unify\_content(Phrase, CCH, CM, DH, DM) to *Priority Unify*

### Rule 3: Priority Unification

TRIGGER: unify\_content(Phrase, CH, CM, DH, DM)  
IF: priority\_unify(CM, CH, C)  
priority\_unify(DM, DH, D)  
THEN: add conceptual\_content(Phrase, C, D) to *Short-term Conceptual Store*

Upper-case identifiers correspond to variables which are instantiated when the rule fires.

Figure 2: COGENT rules for Sense Generation

is of a type of object that has central attributes of *stone* (e.g., inorganic) and has negated central attributes of *lion* (e.g., inorganic, genetic structure not of a lion); it also has diagnostic attributes of *stone* (e.g., a grey colour, hard texture), and some, but not all, of *lion* (e.g., it has a lion shape, but not lion colour or lion habitat).

The case of complex noun phrases involving iterated modifiers is analogous, except that alternate possible syntactic structures (reflected in alternate local parse trees) lead to alternate possible senses. For *fake stone lion*, a sense of *stone lion* may be determined as above. The negating privative *fake* may then operate on this sense (temporarily available in the *Short-term Conceptual Store*) to yield a representation for a type of object that does not have the central attributes of *stone lion* (e.g., inorganic) but does have the diagnostic attributes of this combination. The second reading of *fake stone lion* corresponds to the case where the complex nominal *fake stone* modifies *lion* (corresponding to a lion which is made from some kind of fake stone, such as a hard plastic). In this case, the content of *fake stone* will first be determined. The relevant representation will have as central features the negation of the central features of *stone* (but the diagnostic features of *stone*). The result of combining this with the content of *lion* will depend on the precise form of central features of each. If the central features of *fake stone* do not conflict with those of *lion*, then an affirmative combination will be invoked, and all central attributes will be unified. The resulting sense allows that a fake stone lion is actually a real lion (perhaps pretending to be a statue). Equivocating or negating privative readings may also be licensed if there is conflict on diagnostic or central attributes respectively. In this way, the model demonstrates that the sense generation account of simple privatives can generalise to complex, iterated, concept combinations.

### Future Developments

Future research will focus on four main areas. Firstly, although this work has shown that sense generation can provide an adequate account of iterated privative combinations, the

model may also be used to assess the generality of the sense generation theory in terms of affirmative (i.e., non-privative) forms of combination (e.g., predicating adjectives like *red* and non-predicating adjectives like *criminal*).

The scoping ambiguities inherent in modifier iteration are, of course, just one of a range of syntactic phenomena that falls within the scope of the sense generation theory. A second strand of further work will therefore involve developing the model itself by addressing further syntactic phenomena. For example, combinations may involve syntactic form-class ambiguities. Thus, when interpreting, say, *stone lion*, a strictly bottom-up parser would initially interpret *stone* as a head noun, only to find that it is in fact a modifier. A fuller treatment of syntax would also incorporate the interaction between conceptual content and more coarse-grained content (such as thematic roles) relevant to resolving ambiguity.

A somewhat different area in which the model is underdeveloped concerns the classification and representational functions of concepts. The current model simply takes concept representations to comprise two different types of attribute. The combination operations are sensitive to these types of attributes but not to the particular attributes themselves. This allows a relatively schematic treatment of representation. However, in order to provide a full model of the representational and classification functions, more detail concerning the post-combination specification of the senses is necessary. This is addressed within the sense generation theory (see Franks (1995) for a full account), according to which any such specification is consistent with the results of the initial (schematic) combination stage. These developments will therefore augment (rather than negate) the current model.

Lastly, it is anticipated that developments in the COGENT software will allow further refinement of the model. More detailed modelling of the classification and representational functions of concepts may, for example, be better handled in a hybrid symbolic + connectionist model in which the *Mental Lexicon* (currently modeled as a static, symbolic, store), and its access, is handled by a connectionist component. Affir-

mative combinations in such a model might be handled by a purely connectionist "route", but this route would be interrupted (in the sense of Cooper & Franks, 1993) in the case of privative combinations, where attribute conflicts would trigger type coercion prior to combination.

### Conclusions

We have suggested that a plausible account of concept combination must be constrained by an understanding of the linguistic function of concepts, and that this can begin to be handled by modelling complex, iterated combinations. Given that simple privatives are themselves a complex set of combinations that can provide a test case for a theory of concepts, the ability to model complex privatives provides an even stronger criterion for any account's adequacy. We suggest that the findings reported here indicate that sense generation can provide a framework for developing a plausible, generalisable account of concepts and concept combination.

The development of the model within the COGENT modelling environment has also demonstrated the utility of COGENT as a general modelling resource. We argue that the digrammatic representation (Figure 1), together with the three rules which govern the behaviour of the three processes, clarifies the sense generation theory without obscuring essential aspects with implementation detail. In this sense, we aver, both sense generation and COGENT have profited (and will continue to profit) from their interaction in the development of this model

### References

- Braisby, N. R., Franks, B., & Hampton, J. A. (1994). On the psychological basis for rigid designation. In *Proceedings of the 16<sup>th</sup> Annual Conference of the Cognitive Science Society*, pp. 56–61. Atlanta, GA.
- Braisby, N. R., Franks, B., & Hampton, J. A. (In press). Essentialism, word use, and concepts. *Cognition*.
- Cooper, R. (1995). Towards an object-oriented language for cognitive modeling. In *Proceedings of the 17<sup>th</sup> Annual Conference of the Cognitive Science Society*, pp. 556–561. Pittsburgh, PA.
- Cooper, R. (1996). Perseverative subgoalting in a production system model of problem solving. In *Proceedings of the 18<sup>th</sup> Annual Conference of the Cognitive Science Society*. This volume. San Diego, CA.
- Cooper, R., Fox, J., Farrington, J., & Shallice, T. (In press). A systematic methodology for cognitive modelling. *Artificial Intelligence*, 85.
- Cooper, R., & Franks, B. (1993). Interruptibility as a constraint on hybrid systems. *Minds and Machines*, 3, 73–96.
- Cooper, R., & Shallice, T. (1995). Soar and the case for Unified Theories of Cognition. *Cognition*, 55(2), 115–149.
- Ellis, J., Shallice, T., & Cooper, R. (1996). Memory for, and the organization of, future intentions. In submission.
- Fox, J. (1980). Making decisions under the influence of memory. *Psychological Review*, 87, 190–211.
- Franks, B. (1992). Folk psychology and the ascription of concepts. *Philosophical Psychology*, 5, 369–390.
- Franks, B. (1995). Sense generation: A "Quasi-Classical" approach to concepts and concept combination. *Cognitive Science*, 19, 441–505.
- Franks, B., & Braisby, N. R. (1990). Sense generation, or how to make the mental lexicon flexible. In *Proceedings of the 12<sup>th</sup> Annual Conference of the Cognitive Science Society*. Cambridge, MA.
- Hampton, J. A. (1987). Inheritance of attributes in natural concept conjunctions. *Memory & Cognition*, 15, 55–71.
- Hampton, J. A. (1992). Prototype models of concepts. In van Mechelen, J., Hampton, J. A., Michalski, R., & Theuns, P. (Eds.), *Categories and Concepts: Theoretical Views and Inductive Data Analysis*. Academic Press, London, UK.
- Kamp, H. (1975). Two theories about adjectives. In Keenan, E. L. (Ed.), *Formal Semantics of Natural Language*, pp. 123–155. Cambridge University Press, Cambridge, UK.
- Kaplan, R. M. (1987). Three seductions of computational psycholinguistics. In Whitelock, P. J., McGee-Wood, M., Somers, H. L., Johnson, R., & Bennett, P. (Eds.), *Linguistic Theory and Computer Applications*. Academic Press, London, UK.
- Klein, E., & Sag, I. (1985). Type-driven translation. *Linguistics and Philosophy*, 8, 163–202.
- Medin, D., & Ortony, A. (1989). Psychological essentialism. In Vosniadou, S., & Ortony, A. (Eds.), *Similarity and Analogical Reasoning*. Cambridge University Press, Cambridge, UK.
- Medin, D., & Shoben, E. (1988). Context and structure in conceptual combination. *Cognitive Psychology*, 20, 158–190.
- Miikkulainen, R. (1993). *Subsymbolic Natural Language Processing*. MIT Press, Cambridge, MA.
- Murphy, G. L. (1988). Comprehending complex concepts. *Cognitive Science*, 12, 529–562.
- Newell, A. (1990). *Unified Theories of Cognition*. Harvard University Press, Cambridge, MA.
- Potter, M. C. (1993). Very short term conceptual memory. *Memory & Cognition*, 21, 156–161.
- Pustejovsky, J. (1991). The generative lexicon. *Computational Linguistics*, 17(4), 409–441.
- Rips, L. (1995). The current status of research on concept combination. *Mind & Language*, 10, 72–104.
- Smith, E. E., Osherson, D., Rips, L. J., & Keane, M. (1988). Combining prototypes: A selective modification model. *Cognitive Science*, 12, 485–527.
- Trueswell, J., Tanenhaus, M., & Kello, C. (1993). Verb-specific constraints in sentence processing: Separating effects of lexical preference from garden paths. *Journal of Experimental Psychology*, 19, 528–553.
- Wisniewski, E. J., & Gentner, D. (1991). On the combinatorial semantics of noun pairs: Minor and major adjustments to meaning. In Simpson, G. B. (Ed.), *Understanding Word and Sentence*. North-Holland Publ. Comp., Amsterdam.