

Running head: Dynamic Semantics for Metaphor

Dynamic Semantics for Metaphor

Carl Vogel

Trinity College

University of Dublin

Computer Science

Dublin 2, Ireland

`vogel@tcd.ie`

telephone: 353 1 608 1538

facsimile: 353 1 677 2204

Abstract

An intensional logic with dynamic interpretation is presented in order to provide a formal semantics for sense extension, lexical ambiguity and metaphoricity. Intensionality is required in order to provide the right account of polysemy and homonymy. The dynamics are required in order to allow the interpretation of a sentence to impact the interpretation of subsequent sentences by adding any extended expressions. Metaphoricity is captured in the classification of indices at which expressions are evaluated. A mechanism for deciding which predicates to extend is not provided; the intent is rather to demonstrate how dynamic logic can accommodate sense generation and extension. The system is presented, explained, and argued to capture important features of metaphor creation. It provides existential proof of the potential for formal model theoretic semantics to contribute to the theory of metaphor.

Dynamic Semantics for Metaphor

Introduction

Non-literal language is often thought to be outside the purview of model theoretic semantics. Formal philosophy of language has been influenced by opinions that metaphor, as a form of nonliteral language, is essentially defective or no more than ornamental, even if its use does offer cognitive insights (Percy, 1958). An opposing perspective is that all language use is preconditioned by metaphor, that metaphor is fundamental to cognition and is therefore part of the backdrop to the meaningfulness of sentences rather than something conveyed by them, a view inspired by Lakoff & Johnson (1980). In between is a body of research in artificial intelligence that analyzes metaphoricity through process models (e.g. Fass (1991), Veale & Keane (1992)).

Process models of metaphor interpretation in artificial intelligence research assume that the meaning emerges out of comparisons between domains. Fass (1991), for example, provides a four-way classification of approaches to metaphor: comparison, interaction, selection restriction violation, and convention. However, the comparison approach is present in all of them. Domains are typically encoded as concept hierarchies and comparisons measured via structural morphisms. The system of Veale & Keane (1992), for example, identifies substructures that qualify as the reasonable likeness behind the metaphor. Similarly, Fass (1991) assumes an overarching taxonomy for all domains and discriminates literal, metaphorical and anomalous meanings in terms of relative distance in the hierarchy; that is, structural morphisms are identified within subspaces of the single overarching taxonomy. Indurkha (1994) develops a similar system in which interactionist theories can be explored; in his system, the use of a metaphor can create similarity in domain type hierarchies.

Models of metaphor understanding devoted to the process of identifying structural preconditions for metaphor do not provide an entire theory of interpretation for nonliteral language: such a system conveys much about what a metaphorical sentence could mean, but does not offer insight into whether the sentence is true or not. This paper focuses on the truth conditions of metaphor and the integration of metaphorical expressions with a standard formal framework for the syntax/semantics interface. The paper does not provide an alternative account of what counts as a good metaphor, nor the relevant analogies or lack of alternative lexical items that gives rise to new metaphors. I assume that a suitable account of such can be borrowed from the extensive literature. It is an assumption of this paper that meanings must be delivered for metaphorical sentences using the same formal apparatus as the literal senses, albeit with the locus of metaphoricity appropriately identified within the system. A major point is to demonstrate that metaphoricity is not outside the remit of natural language semantics. Rather, explanations of certain aspects of metaphor are integral to the theory of meaning of any sentence in a natural language.

Truth conditions are a small part of meaning, but a profoundly essential part. Without truth conditions, comparison cannot happen: in order to compare two sets it is essential to be clear on what comprises the membership criteria for the categories independently, even if the criteria are vague or ill-defined for some compared sets (these criteria are instances of truth conditions). In general, comparing two entities is parasitic on being able to individuate the compared entities, and the degree to which individuation is possible is highly correlated with the capacity to give truth conditions to a sentence asserting an individuation property. This is consistent with psycholinguistic models such as that of Glucksberg & Keysar (1993) which analyze metaphors as categorization statements; there as well everything hinges on the extent to which entities are contained in categories. In overlooking truth conditions, purely structural theories are unable to characterize certain dynamic properties of metaphoricity, e.g. that interpreting a

metaphor can change the interpreter's concept of the world (note that Indurkha (1994) does not overlook truth conditions and also recognizes dynamic aspects). Of course, when an interpreter accepts the veracity of a declarative assertion, a sort of change of world is brought about, but this is orthogonal to the kind of change that happens with metaphor. Metaphor brings about a change of possible concepts. It is at the heart of the ontogenesis of literal language. A full theory of metaphor in semantics requires an account of truth conditions inclusive of a theory of the impact on subsequent interpretation.

It may seem that the truth conditions of metaphor are trivial. Metaphors are simply literally false (or, when negated, patently true), while their counterparts expressed as similes cannot be false. This is well-known, as is the additional fact that a metaphorical assertion can be true or false in its own nonliteral terms.

- (1) Leslie is a library.
- (2) Leslie is like a library.

From these issues with truth-values, some have concluded from this that a semantic theory for natural language that relies on a notion of truth will have a hard time articulating a theory of meaning for metaphorical sentences. Davidson (1984), in fact, argues that metaphoricity is indeed a property of language use, and hence not the business of semanticists. Certainly, he is not alone in that view (Morgan, 1993).

However, Vogel (1998) has argued against the pessimistic extreme of this view and demonstrated that certain aspects of the pragmatics can be captured in a straightforward model theoretic account, and that account is substantially extended here. van Genabith (this volume) has provided a type theoretic treatment which analyzes metaphors as reduced similes, with an analysis that differs significantly from a related formal account proposed by Miller (1993). A difference with my proposal is that I agree with Davidson (1984) that metaphor should not be analyzed via translation to simile. The truth

conditions differ, and there is not a guarantee of a unique simile to translate a metaphor into, and from which to elide instances of “like”, for example in the nonliteral use of any verb other than the copula. van Genabith claims that the truth value complementarity between metaphor and simile is mitigated when “trivial” likeness relations are ignored. This requires, in turn, that the simile (2) be translated into the reduced typicality sentence (3)

(3) Leslie has a property that is a typical property of libraries.

Actually, typicality isn’t essential to his account, but it is the best motivated choice of nontrivial properties to assume. In fact, any nontrivial property will do. The reductionist account by its essential nature omits an important property of metaphor: similes, even restricted to existentially quantifying over interesting properties, do not have the “special” force that metaphors do. Glucksberg & Keysar (1993) argue that metaphors are generally perceived as stronger than related similes. They also notice that a sentence needn’t have a unique related simile to translate a metaphor into. While it is technically possible to say that the property shared by tenor and vehicle is the special metaphorical one, such a move undermines the intuitive appeal of the type-theory implementation of the comparison approach to a metaphor as if an existential assertion of nontrivial likeness of named categories. Maintaining nontriviality does allow metaphorical statements to be contingent, even on the nonliteral interpretation, but the part of the analysis that captures the nonliteral interpretation does so wrongly in my opinion by equating the metaphor meaning with that of a reduced simile. Metaphors involve (and their first uses create) special senses of the expressions at stake.

Vogel (1998) gave a first order logical language (see Partee, Meulen, and Wall (1993) for an accessible presentation of foundations of first order logic) in which literal and non-literal utterances can be expressed and discriminated. As a first order account, it allows variables over individuals, but not over relations between individuals (the type logic

account of van Genabith (this volume) is one with variables over relations, and hence is not a first-order system). It is extensional in assuming that any predication is exhaustively specified by providing the set of individuals that stand in the named relation. This means that in a very basic formal system, one which is completely extensional in its analysis of meaning (in that the meaning of a term is fully specified by the set of items that the term truthfully denotes), it is possible to provide an account of metaphoricity in natural language. An advantage of a logical approach such as the one proposed here (or of van Genabith (this volume)) is in its methodology: we understand completely the syntax and semantics of the language, and therefore we can be fully explicit in stating the theory of metaphor in its terms, as well as how the theory of metaphor integrates with other semantic phenomena.

There are two main ways in which the system of Vogel (1998) diverges from classical uses of first order logic as a language for meaning representation. Firstly, models for the language initialize each predicate in the language with two characteristic sets rather than one as is usually the case. One of the characteristic sets is the set of objects that satisfy the predicate literally, and the other set, initially empty, is the set of objects which satisfy the predicate nonliterally. Secondly, the approach adopts techniques from dynamic semantics (e.g., Groenendijk & Stokhof, 1991). The interpretation of sentences has a dynamic impact on the models. Essentially, certain nonliteral expressions have the capacity to add elements to the characteristic sets of predicates involved in the metaphorical sentence under interpretation. The output of the interpretation of one sentence is the input to the interpretation of subsequent sentences.

This approach correctly discriminates the truth conditions of metaphors and similes without handling them as reduced similes, yet creating a concrete reason for a related explicit comparison statement to be true. That is, the extensional unpacking of ‘being the same type as’ is for the categories to have an element in common. The account

provided by Vogel (1998) allows that a successful metaphor constructs a situation in which two categories do have a common element, and thereafter the comparison sentence is also true. In this framework, the metaphorical expression remains literally false, although it is true with respect to the nonliteral interpretation. Moreover, the approach accommodates the dynamic aspect of meaning in such nonliteral language—interpreting a nonliteral sentence extends the meaning of predicates at issue by adding nonliterally predicated entities to the corresponding characteristic sets. Thus, the framework provides an analysis of sense extension at the same time. Vogel (1998) discusses certain syntactic constraints that seem to be in place to allow or prevent sense extensions from occurring.

However, this model is not rich enough to make all of the required discriminations. In particular, the approach does not allow for there to be more than one way for a predicate to be used nonliterally: assuming that (1) was intended nonliterally, is Leslie being described as knowledgeable or as a lender? This failing is a consequence of the framework's inability to deal with lexical ambiguity in general, even for literal predicates. This paper addresses these problems by recasting the main dynamic interpretation ideas from Vogel (1998) in a more expressive intensional setting. Rather than just a first order logic, the approach proposed here is a modal predicate logic (see Hughes & Cresswell (1968) for an overview).

Modal predicate logic derives from the observation that not all logical expressions are truth functional. The specific modalities examined can vary, but typically there is a 'box' operator and a 'diamond' operator. The box operator involves universal quantification over alternatives, and the diamond operator involves existential quantification over alternatives. Thus, necessity would be modeled with the box, and possibility with the diamond, and similarly for other dual modal relations, epistemic, deontic, etc. The operators are not truth functional in that "possibly p " doesn't depend for its truth just on the value of p in the real world, but also on alternatives. So, "possibly

p ” is true if and only if p is true in some accessible alternative, even if p is false in actuality. Axioms appropriate to the modality at stake have underneath them constraints on the accessibility relations that connect possible alternatives. The alternatives themselves are modeled as possible worlds, propositions which may be either partial or total in the sense of each predicate being given as true or false in each alternative. At each world we have a set of things, and relations among them, just as in a first order account. Each world is a way things might have been. It is intensional in acknowledging that predicates are not fully determined by their extensions in the actual world; rather, to understand a predicate is to know its extension, the set of elements it is true of, in each alternative.

In the current approach as well, for each predicate, a different possible world provides the characteristic set corresponding to the particular sense at stake. It is not that an entire world is metaphorical or literal, but that a particular world counts as metaphorical for a predicate because of the entities that comprise the predicate at the world. Then, the contrast between literal and nonliteral meaning is based upon the predicate-relative classification of the worlds themselves. A world may be literal relative to some predicate and nonliteral for another. This offers an elegant way of modeling the meaning shift that occurs when a metaphor dies: the sense doesn’t change, but the classification of the sense as literal or nonliteral does. Thus, the interpretation mechanisms for metaphorical expressions are exactly the same as for literal expressions, and estimations of degree of metaphoricity are external classifications of the senses used in the interpretation. The classical modal operators correspond in this framework to quantification over senses: $\Box p(q)$ means that $p(q)$ is true (q is p) no matter which of p ’s senses is considered, and $\Diamond p(q)$ means there is some sense of p for which p is true of q . Additionally, we can have modal assertions that make explicit reference to a particular sense, rather than using quantification (e.g. “Leslie is a fox in the cleverness sense of the term”).

A greater range of potential metaphorical expressions are handled by the current proposal, individual terms (e.g., names), in addition to predications that take terms as arguments. The system of Vogel (1998) could handle metaphorical uses of (1) and (4).

(4) Dr. Smith hit her patient with bad news.

(5) Einstein here [speaker points] says he knows how to start the grill.

However, this system did not have a convenient way of handling (5), as it did not permit sense extension for constant terms.

Related ideas are discussed by Hintikka & Sandu (1994). They notice as well that the possible worlds semantics for modal logic can be applied to the theory of metaphor and polysemy. They refer to a “meaning line” across possible worlds. This is related to the analysis of identity of individuals in possible worlds semantics for first-order logic. It is standard to assume that at each possible world it is possible to identify counterparts of individuals at other worlds. A meaning line gives not a counterpart relation for individuals, but for predicate names. The meaning line across worlds indicates what entities are in the characteristic set of individuals at each world. The current framework can be seen as giving more concrete detail to a proposal of Hintikka & Sandu (1994), while also drawing out further aspects of metaphoricity and polysemy that can be captured in the framework (notably, sense extension) by adding dynamics.

An Intensional Dynamic Semantics for Sense Extension

The formal presentation of the account is relegated to the appendix. Here I will discuss the main features.

Syntax of the Language

The account presumes that it is possible (but not necessary) to indicate the sense in which an ambiguous expression is intended. It is possible further to indicate whether an

expression is used literally or metaphorically. “Body language” accompanying an utterance can be used for this. Goatly (1997) provides a more exhaustive litany of metaphoricity cues (including explicit use of markers like “metaphorically speaking” and even “literally”, ironically enough). It happens that the cues, (apart from “metaphorically”) can be used for other purposes as well; hence, they tend to be ambiguous. However each can be interpreted by a listener as signaling some sense or other, even if not actually signaling one such. Thus, natural language includes more and less explicit designations of sense. Interpretation, in absence of a signal, is relative to the sense a hearer finds germane.

I assume that senses can be given a partial order. One possible ordering is the frequency relation—the frequency with which the term is used with a particular sense. Another more appropriate to the topic at hand is degree of ‘liveliness’ of a sense, where liveliness corresponds to metaphoricity. Goatly (1997), for example, gives a 5-way classification of degrees of conventionalization of metaphor. Death of metaphor is a transformation into literalness, and in the current system this is modeled with rearrangements of the partial ordering of senses.

Designations of sense are quite like the modal operators. Instead of universal or existential quantification over senses, they admit the possibility of referring to a particular way things can be directly. Similar mechanisms are used in modal construals of tense logic (e.g. “On Wednesday Morning, 3 a.m., GMT, November 11, 1965....”); there the alternative worlds are moments in time rather than possible predicate meanings. Indications of sense are assumed to be iterable (also like tense operators—“...did want to become...” vs. “will want to have been...”) I will assume that as with tense, the outermost (in English, the leftmost) designation sets the ultimate reference point.

- (6) In Freud’s sense of Marx’s sense of ‘repression’, economic exploitation is the result of frustrated desires from childhood.

- (7) In Marx's sense of Freud's sense of 'repression', frustrated desires from childhood are the result of ownership construals of personal relations
- (8) Freud's sense of Marx's sense of 'repression' is the same as/different from Marx's.

In each case, the designating term is the highest or outermost one in syntactic terms.

In addition to referring to senses, it is also assumed possible to refer deictically. Thus, the language also 'includes' deixis. This is modeled by a function that maps pointing acts to elements of the domain. Deictic acts fix reference, and in idealization I will assume that deictic acts are unambiguous. Deixis is also assumed to accompany other referential expressions. It would be possible to allow deictic acts to stand on their own, satisfying argument roles of verbs without other expressions. But that generality would reduce expository clarity. Thus, deictic acts are seen here as functioning like resumptive pronouns in failing to reduce predicate valency, but distinctly in allowing nonresumptive reference. Like designations of sense, deictic acts are possible in the language, but not necessary. Deixis makes it easy to extend the sense of a name, but is not essential there, nor in extending other nonlogical constants.

A separate axis of interpretation of the language, akin to a listener determining or ascertaining a signal for a sense, is assertional. An utterance can be deemed by an interpreter as new information to be added to world knowledge, or as disputable information. New information, the interpreter will accept as an utterance and adapt conceptualization to. Disputed information is that retrieved from utterances not accepted as being new and true. Disputable utterances are tested as true or false. I will not consider information retraction here. The category an utterance falls into along this dimension is patently determined by pragmatic factors. In the present analysis, this boils down to how the listener chooses to interpret an utterance. Taking a sentence as assertional involves evaluating it with a dynamic interpretation function; understanding a sentence as up for

debate involves evaluating it statically. Any predication, metaphorical or literal, can be a constituent of an utterance that causes a listener to modify beliefs or which simply functions as a test. It is up to the interpreter to decide on how to interpret, on the basis of any explicit cues attended to by the listener or simply on the basis of proclivities.

Interpreting the Language

The formal interpretative rules are supplied in the appendix. The basics out of which the rest is constructed are as follows: there is a set of indices corresponding to possible worlds (possible senses of predicates). There is also a fixed domain, and an interpretation function which maps individual constants (names) onto elements of the domain, and predicates of arity n onto n -place relations constructed from the domain. Each index has a unique identifier. Designations of sense are mapped uniquely to indices. Each predicate has only one characteristic set at an index, its extension at that index. Deictic acts are interpreted by a function that maps them directly to an element of the domain, much in the way assignment functions interpret variables. Assume that the entire domain is available at each index.

Interpretation of an utterance has an input and an output. The output of interpretation of one utterance is input to the interpretation of subsequent utterances. There are a number of possible ways for a listener to interpret an utterance, depending on whether the listener takes the information as new or disputable. The input and output to interpretation is the domain and interpretation function mentioned above. Assuming the domain is constant for simplicity, what can change over the course of interpretation is precisely the interpretation function—what constants point to what in the domain at an index; what extension a predicate has at an index; what indices exist.

Take the case of a literal unambiguous expression. It is true at an index if the index is among the literal ones and the entities referred to stand in the mentioned relation.

Exactly the same analysis holds for metaphorical expressions, except that the necessary indices come from a different class of indices. In fact, the only difference between literal truth and metaphorical truth is whether the relevant senses happen to be classified as such. The forms of evaluation just described are both static. These are used to test the truth values of potentially disputable information. Neither changes the overall state of information. The test of truth amounts to set membership of some entities in the extension of a predicate. The degree of metaphoricity of the sense is a separate issue.

The other possibility is dynamic interpretation. Here also a sense may be indicated, but using dynamic interpretation a new one may be generated, as in the case of absolutely novel metaphor. In this case, the listener has decided to accept the information update supplied by a sentence (literal or nonliteral). The result is simply that a designated individual (in the case of a constant or a unary relation) or a tuple in the general case, is added to the interpretation function. In case the designated index does not yet exist, we have a novel metaphor. In other cases, we have extension of existing senses, literal or nonliteral.

In static interpretation, the input function is the output function. In dynamic interpretation the output function can be distinct (here I consider on possibilities of monotonic increase; however, of course, contraction is an important topic in the belief revision literature). Dynamic interpretation can involve the creation of new indices or extending characteristic sets of predicates at existing indices. Take the later case first as it is common to both literal and nonliteral sense extension: simply, the characteristic set of the designated predicate at an index is extended to include additional elements. In the case of generating a new sense altogether, the world given as the input to interpretation is taken as the standard—all the denotations of other predicates unrelated to the extended predicate maintain their existing characteristic sets. The extended predicate and any related predicates are stipulated as having in their extension the focussed tuple. The

result is available for subsequent discourse. The theory does not offer a method for deciding which other predicates to extend. Nor does it stipulate a method for identifying which world to extend when the sense is not signalled. Rather, the system is compatible with total ambiguity. It would be quite useful to explore what theoretical approaches to word sense disambiguation dovetail with the proposed approach most naturally.

The framework does not offer an explanation of when to adopt a dynamic interpretation as opposed to a static one. Dynamic interpretation nearly always succeeds when adopted (and even in cases in which the tests involved in static interpretation would have yielded a ‘false’), and sometimes causes a change to the interpretation function. Importantly, static interpretation is not a necessary initial step in the understanding of a metaphor. It is not necessary to interpret a sentence literally first and then metaphorically.

The framework does offer a way to explore the syntax/semantics interface for sense extension of literal and nonliteral expressions. I personally have a hunch that negation is static. That is to say, I believe that the sense of an expression cannot be extended within the scope of negation. This is a very different claim from the possibility of a metaphor being used within the scope of negation.

(9) Leslie is not a newt

(10) Leslie is not an accountant

It seems that new senses cannot be generated under the scope of negation, and further that existing senses cannot be extended in the same environment. This and related constraints can be modeled in the system by stipulating static interpretation for certain logical (and perhaps nonlogical) constants. It is an advantage of the approach that it affords room for such explorations.

Discussion

The interpretation function for nonliteral expressions in general creates a situation in which the expression is nonliterally true, regardless of literal truth values. This system releases the requirement of Vogel (1998) that a prerequisite to extension be the literal falsity of the expression. This means that (1) can be simultaneously literally and nonliterally true (provided, for example, that Leslie is literally a library, and also a library in the designated special sense). Interpreting a nonliteral expression with respect to a static interpretation function allows nonliteral expressions to be true or false, with respect to whichever index happens to be the default or signaled index. Sense extension cannot happen using a static interpretation function, but re-use of an extended sense can: once created a metaphor can be re-used as if literal. This correctly captures the fact that (1) can be false, even when used metaphorically, if it isn't the case that Leslie is in the characteristic set for the special sense of "library". The system also correctly analyzes the first use of a metaphor (subject to certain syntactic restrictions and interpretational decisions) as inescapably (but nonliterally) true.

The difference between literal and nonliteral in this system is not equated with interpretation using a fully static interpretation versus using a dynamic interpretation (see definition 3 in the Appendix). Rather it is in the classification of the index at which (either static or dynamic) interpretation occurs. This means that there is no commitment to a strong division between literal and figurative language. There are only senses of terms, and classifications of senses. It is up to the individual interpreter to decide what the relative figurativeness ordering is. This classification consists here in grouping the index at which the interpretation for an expression occurs as within \mathcal{N} , the set of indices that are nonliteral for a word, or \mathcal{L} , the set of indices that are literal for a term. Indices can also be partially ordered to account for the fact that the literal/nonliteral distinction is not binary. The death of metaphor in this system involves nothing else but the

reclassification of the index at which the metaphorical expression is interpreted as being no longer within \mathcal{N} but \mathcal{L} .

In connection with these points, an anonymous reviewer has pointed out that some sentences may strike some people as metaphorical and others as completely literal. The reviewer supplied the following example which demonstrates the interaction of graded categorization with perceptions of metaphoricity.

(11) Olive oil is nature's more versatile fruit juice.

I agree completely that the literal/figurative divide is arbitrary, and I hope to have captured that arbitrariness in the model in the simple act of classification of the senses. Any one interpreter will make certain decisions about what counts as a metaphorical sense and what doesn't. In particular, that senses can be partially ordered by the interpreter captures the graduation of perceived metaphoricity.

Just as the accessibility relations in classical applications of modal logic provide interpretations for modal axioms, here the accessibility relations capture similarities of sense. That is, work on conceptual metaphor can be related to the current proposals by noticing that conceptual metaphors, when expressed in particular sentences, tend to correlate senses of the terms used. It will be interesting to explore the exact properties that different forms of accessibility on sense indexes will have in the interpretation of meaning relations among terms.

Given some expression to interpret at a nonliteral index that extends the characteristic function for the expression at the index, there are likely to be other expressions that also require extension. These ancillary expressions are related to the first through an initial theory of the world.

(12) Leslie is a fox.

(13) $\forall x \text{ Fox}(x) \rightarrow \text{Mammal}(x)$

(14) $\forall x \text{ Fox}(x) \rightarrow \text{LikesToStealChickens}(x)$

Given the constituent expressions of (12), it is reasonable to imagine that they participate in other sentences that constitute an interpreter’s theory of the world, as in (13–14).

However, imagine the first nonliteral use of (12) in which “fox” designates a sense corresponding to that of “sly”. The denotation for “fox” at the new sense is structured so that it has Leslie within it. However, it is also reasonable to consider under this same sense whether other predicates connected to “fox” also require extension to cover Leslie in this new sense. Assuming that Leslie is a human being, a plausible system for deciding which other predicates to extend would leave “mammal” untouched (as Leslie is already there literally), but might be inclined to add Leslie to the set of things who like to steal chickens in a nonliteral sense that corresponds to the nonliteral sense of “fox” under consideration. Identifying which predicates are pertinent to a metaphor’s implicative complex is exactly the business of structural mapping process models. I simply assume that one or other of them (e.g. Veale & Keane, 1992) can stipulate which additional predicates need to be extended and which ones to leave alone. The current system expects that the other module would deliver in turn each predicate requiring extension, as definition 3 requires (see Appendix) in general that each extended expression be atomic. Coordination and universal quantification are exceptions, but those just reduce recursively to atomic extensions.

Conclusions

Nonliteral expressions can be nonliteral because of the predicate or because of the argument. The proposed semantics allows extension from the meaning of both basic predicates and arguments. Extension of the interpretation of a constant can fail if there isn’t an accompanying deictic act to make clear what the constant is to be extended to, but if deixis is present, a constant can have an extended sense whether used in the scope

of a literal predicate or a nonliteral predicate. When used in the scope of a literal predicate, the predication can turn out to be false. However, when used in the scope of a nonliteral predicate, the sentence will evaluate as true, and will extend the appropriate sense accordingly. It remains possible to extend each sense to the point of triviality by applying it to all elements of the domain, but if everything is spoken nonliterally then discriminations of meaning are meaningless. The meaning of a metaphorical sentence is not reduced to the meaning of a simile—it creates a reason for a related simile to be true: nonliteral expressions create an intersection of denotations which thereby licenses a nontrivial likeness between predications involved. Reclassification of indices as members of \mathcal{N} or \mathcal{L} captures the drift of new metaphorical senses to literalness.

Appendix: An Intensional Logic for Sense Extension, Polysemy and Metaphor

Syntax

Meanings of sentences are represented here by translation into a language that allows (but does not require) exact specification of the sense of a predicate. Explicitness about which is the intended sense also conveys the classification of the sense as literal or nonliteral. There can be a signal of the sense that is used, and if it is present it may be accordingly interpreted; if it isn't present then interpretation will be made relative to a sense determined otherwise. The syntax admits, for example, c_{m_m} etc. as constants. The potential iteration survives in the semantics.

Dfn. 1 *Syntax*

1. Assume a set of constants, C , a supply of variables V , predicates \mathcal{R} , indications of sense M , and the usual connectives.

2. If c is a constant and m is an indication of sense, then c_m is a constant.

3. A constant may be accompanied by a deictic act.

4. *If P is an n -ary predicate name, $n > 0$, and m is an indication of sense then P_m is a predicate name.*
5. *If P is well-formed, then so is $\forall xP$*
6. *If P is well-formed, then so is $\neg P$*
7. *The usual combination rules with respect to forming predications, complex formulae and sentences (truth denoting expressions with no unbound variables) apply*

The language provides a basic intensional system that lacks the usual primary focus on interpretations for quantificational modalities (modal operators have interesting interpretations, but this isn't the issue here). A predicate has a characteristic set at each 'world'. A predicate may be used in a way that indicates at which world it should be evaluated. The main modal operators of interest are ones that select an explicit index rather than quantifying over indices. It is further possible for a complex sentence to select one world for one predicate and a different world for a different predicate in the compound. If no sense is indicated for a predicate, interpretation is relative to a particular world. The same is true for constants. "World" and "index" are synonymous mathematical terms here; they carry no ontological baggage. Assume that the same domains are available at each index, but let the interpretation of constants fluctuate. This will account for the interpretations of sentences like (1) that use the copula as well as more complicated predications like (4) and (5).

Semantics

Let D be a nonempty domain and W be a set of indices. The interpretation function I for basic expressions in the language is presented in terms of the tuples comprising it. Assignment functions g map variables to elements the domain; sense selection functions s map sense indicators to indices; and deixis functions δ map pointing acts to elements of the domain. Let \mathcal{L} be a subset of W corresponding to the indices for

literal senses, and define \mathcal{N} as $W - \mathcal{L}$

Dfn. 2 *the basic interpretation function, I*

1. $\forall c \in C, w \in \mathcal{L}$ there is a unique $d \in D : \langle c, w, d \rangle \in I$.
2. $\forall P^n \in \mathcal{R}, n \geq 0, \forall \tau \in D^n, \langle P, w \rangle \oplus \tau \in I$ iff P is true of the tuple τ at index w .

This specifies the initial interpretation function. When the arity of the predicate (given by a subscript; below arity is indicated only on basic predications, not on possibly compound formulae) is 0, τ is the empty tuple, and P is just a proposition. The symbol \oplus denotes sequence concatenation. The meaning function ($\llbracket \cdot \rrbracket$) for arbitrary expressions in L is defined below. This function depends on the interpretation function I , variable assignment g , sense indication s , an index w , and deixis δ . Static or dynamic versions can be called upon at indices given by w or s whether the index is literal or nonliteral. Roughly, static interpretation is a classical special case of the dynamic interpretation clauses given here. Interpretation functions appear to the left and to the right of the meaning function. When interpretation is static, it is exactly the same interpretation function on both sides. With dynamic interpretation, the right hand side has additional tuples, and the resulting interpretation function becomes the input.

Dfn. 3 *dynamic interpretation*

1. $I \llbracket c \rrbracket^{I \cup \{ \langle c, w, \delta(c) \rangle \}, \langle s, g, \delta, w \rangle} = \delta(c)$, iff $\delta(c)$ is defined
2. $I \llbracket c \rrbracket^{I, \langle s, g, \delta, w \rangle} = I(c, w)$, iff $\delta(c)$ is not defined but $I(c, w)$ is, and is otherwise undefined
3. $I \llbracket c_m \rrbracket^{I \cup \{ \langle c, s(m), \delta(c) \rangle \}, \langle s, g, \delta, w \rangle} = \delta(c)$, iff $\delta(c)$ is defined
4. $I \llbracket c_m \rrbracket^{I, \langle s, g, \delta, w \rangle} = I(c, s(m))$, iff $\delta(c)$ is not defined but $I(c, s(m))$ is, and is otherwise undefined

5. $I[[x]]^{I, \langle s, g, \delta, w \rangle} = g(x), \forall x \in L$
6. $I[[t^1, \dots, t^n]]^{O, \langle s, g, \delta, w \rangle} = \langle I[[t^1]]^{O^1, \langle s, g, \delta, w \rangle}, \dots, O^{n-1}[[t^n]]^{O, \langle s, g, \delta, w \rangle} \rangle,$
7. $I[[P^0]]^{I, \langle s, g, \delta, w \rangle} = 1$ iff $\langle P, w \rangle \in I$
8. $I[[P^n(\sigma)]]^{I \cup \{ \langle P, w \rangle \oplus I[[\sigma]]^{O, \langle s, g, \delta, w \rangle} \} \cup O, \langle s, g, \delta, w \rangle} = 1$ iff $n > 0, |\sigma| = n$
9. $I[[P_m^n(\sigma)]]^{I \cup \{ \langle P, s(m) \rangle \oplus I[[\sigma]]^{O, \langle s, g, \delta, w \rangle} \} \cup O, \langle s, g, \delta, w \rangle} = 1$ iff $n > 0, |\sigma| = n$
10. $I[[\neg P]]^{I, \langle s, g, \delta, w \rangle} = 1$ iff $I[[P]]^{I, \langle s, g, \delta, w \rangle} = \emptyset$
11. $I[[P \wedge Q]]^{O, \langle s, g, \delta, w \rangle} = 1$ iff $I[[P]]^{M, \langle s, g, \delta, w \rangle} = 1$ and $M[[Q]]^{O, \langle s, g, \delta, w \rangle} = 1$
12. $I[[\forall x \phi]]^{O^U, \langle s, g, \delta, w \rangle} = 1$ iff $O^U = \bigcup O^i$, where $I[[\phi]]^{O^i, \langle s, g[x/d], \delta, w \rangle} = 1,$
 $\forall d \in D.$

Output interpretation functions (e.g., O^i and M) are the smallest ones satisfying the conditions.

Definition 3.1-4 handle nonliteral constants: either it is a first and deictic use, or it is a reuse of a previously extended sense. The first use is the case of sense extension—the interpretation function that is the output of interpretation has an additional tuple in it whose content depends on deictic reference (3.1 and 3.3). This is at the heart of what was referred to earlier when mentioning that interpreting a new metaphor changes the interpreter’s concept of the world—a new sense exists with referents, and this sense is available to interpretation of subsequent discourse. Subsequent use of the same extended sense may apply it to the same individuals (3.1-4) or to additional elements of the domain (3.1 and 3.3). In the case that deixis isn’t given, then it can only be a subsequent use of the sense-extended constant, and is thus static with respect to the interpretation function, whether or not the sense is indicated (3.2 vs. 3.4). The interpretation of variables is unaffected by polysemy (3.5). The interpretation of a sequence of constants and variables

is sequential, and the output of the interpretation of each term in the sequence is the input to the interpretation of the next one (3.6); this echoes the widely assumed asymmetry of argument structure in natural languages. Senses may be signaled or not for any of the terms, as in the absence of a sense indication, interpretation is relative to the designated index w . Propositions are not given a dynamic interpretation (3.7). Definition 3.8 and 3.9 allow the sense of a predicate name to be extended. In 3.8 the sense is extended at a nonliteral index independently specified, and in 3.9 it is relative to an indicated sense. Note that the output interpretation function includes additions made in the interpretation of the predicate's argument. Negation (3.10) has a static interpretation, but conjunction is dynamic (3.11)—the second conjunct of an expression can be interpreted relative to the extended interpretation from the first conjunct. This means that the second conjunct can be a literal predication of arguments with a nonliteral meaning created by prior discourse. Finally, universal quantification (3.12) extends meaning by iterating over all elements in the domain in combinations of the preceding possible ways of extending meaning. That the intermediate interpretation functions in clauses (3.11) and (3.12) are the minimal ones that work means that arbitrary choices will not do, only recursively constructed extensions of the initial interpretation. This definition of the meaning function, $\llbracket \cdot \rrbracket$, is just one of the possible specifications an interpreter can make use of. Another is fully static.

Acknowledgements

I am grateful for enlightening debate on these topics with Josef van Genabith and Tony Veale, also for extremely useful feedback from various reviewers and editors, especially John Barnden. This research is supported in part by Forbairt Basic Research Grant SC-97-623.

References

- Davidson, D. (1984). What metaphors mean. In D. Davidson (Ed.), Inquiries into truth and interpretation (p. 245-64). Oxford: Oxford University Press.
- Fass, D. (1991). met*: A method for discriminating metonymy and metaphor by computer. Computational Linguistics, 17(1), 49-90.
- Glucksberg, S., & Keysar, B. (1993). How metaphors work. In A. Ortony (Ed.), Metaphor and thought (2nd ed., p. 357-400). Cambridge University Press. (First Published, 1979)
- Goatly, A. (1997). The language of metaphors. Routledge.
- Groenendijk, J., & Stokhof, M. (1991). Dynamic predicate logic. Linguistics and Philosophy, 14, 39-100.
- Hintikka, J., & Sandu, G. (1994). Metaphor and other kinds of nonliteral meaning. In J. Hintikka (Ed.), Aspects of metaphor (pp. 151–87). Kluwer Academic Publishers.
- Hughes, G. E., & Cresswell, M. (1985). An introduction to modal logic. Methuen and Co.
- Indurkha, B. (1994). Metaphor as change of representation: An interaction theory of cognition and metaphor. In J. Hintikka (Ed.), Aspects of metaphor (p. 95-150). Kluwer Academic Publishers.
- Lakoff, G., & Johnson, M. (1980). Metaphors we live by. Chicago, IL: Chicago University Press.
- Miller, G. (1993). Images and models, similes and metaphors. In A. Ortony (Ed.), Metaphor and thought (2nd ed., p. 357-400). Cambridge University Press. (First Published, 1979)

Morgan, J. (1993). Observations on the pragmatics of metaphors. In A. Ortony (Ed.), Metaphor and thought (2nd ed., p. 124-34). Cambridge University Press. (First Published, 1979)

Partee, B., Meulen, A. ter, & Wall, R. (1993). Mathematical methods in linguistics. Kluwer Academic Publishers.

Percy, W. (1958). Metaphor as mistake. Sewanee Review, 66, 79-99.

van Genabith, J. (2000). Metaphor, logic and type theory. (this volume)

Veale, T., & Keane, M. (1992). Conceptual scaffolding: A spatially founded meaning representation for metaphor comprehension. Computational Intelligence, 8(3), 494-519.

Vogel, C. (1998). A dynamic semantics for novel metaphor. In Proceedings of the third international conference on information theoretic methods in logic, language and computation (p. 116-27).