Gentner, D., Bowdle, B., Wolff, P., & Boronat, C. (2001). Metaphor is like analogy. In Gentner, D., Holyoak, K.J., & Kokinov, B.N. (Eds.), *The analogical mind: Perspectives from cognitive science* (pp. 199-253). Cambridge MA, MIT Press.

6

Metaphor Is Like Analogy

Dedre Gentner, Brian F. Bowdle, Phillip Wolff, and Consuelo Boronat

Metaphor is pervasive in language and thought: in scientific discovery (Gentner 1982; Gentner and Jeziorski 1993; Gruber 1995; Nersessian 1992), in literature (Gibbs 1994; Lakoff and Turner 1989; Miller 1993; Steen 1989) and in everyday language (Glucksberg and Keysar 1990; Hobbs 1979; Lakoff and Johnson 1980). Yet despite a considerable amount of research, surprisingly little is known about how metaphors are psychologically processed.

In this chapter we present an approach that unifies metaphor with processes of analogy and similarity. We first lay out the analogy approach to metaphor and delineate some limitations. In the first section we ask whether large-scale conceptual metaphors such as *Love is a journey* or *The mind is a computer* can be modeled as extended structural mappings between domains. Our research suggests that the answer is *yes* for novel metaphors, but not for conventional metaphors. In the second section we describe research that shows that the real-time processing of nominal metaphors can be captured by detailed models from analogy. In the third section we lay out a theory—the *career of metaphor*—of how metaphoric representation changes as a metaphor evolves from novel to conventional. In the fourth section we discuss implications of these ideas.

One reason metaphor[1] is challenging is its range of types, as in the following list:

1. A man is not necessarily intelligent because he has plenty of ideas, any more than he is a good general because he has plenty of soldiers (Chamfort).
2. My job is a jail.

3. His eyes were burning coals.
4. Tires are like shoes.
5. On a star of faith pure as the drifting bread, /As the food and flames of the snow (Dylan Thomas).

Metaphors (1) and (2) could be considered analogies—comparisons that share primarily relational information. But metaphors can also be based on common object attributes, as in (3), or both, as in (4). Most of the metaphors studied in the psychological literature are analogies—that is, they convey chiefly relational commonalities: for example, *Encyclopedias are gold mines, My job is a jail.* According to structure-mapping theory, such relational metaphors convey that a system of relations holding among the base objects also holds among the target objects, regardless of whether or not the objects themselves are intrinsically similar. The centrality of relations during metaphor comprehension has been confirmed by a number of studies. For example, people's interpretations of metaphors tend to include more relations than simple attributes, even for statements that suggest both types of commonalities (e.g., Gentner and Clement 1988; Shen 1992; Tourangeau and Rips 1991). Further, Gentner and Clement (1988) found that subjects' judgments of the aptness of metaphors was positively correlated with the judged relationality of their interpretations of these metaphors, but was either negatively correlated or unrelated to the attributionality of their interpretations.

According to structure-mapping theory (Gentner 1983, 1988; Gentner and Markman 1997), analogical mapping is a process of establishing a *structural alignment* between two represented situations and then projecting inferences. Structure-mapping theory assumes the existence of structured representations made up of objects and their properties, relations between objects, and higher-order relations between relations. An alignment consists of an explicit set of correspondences between the representational elements of the two situations. The alignment is determined according to *structural consistency* constraints: (1) one-to-one correspondence between the mapped elements in the base and target, and (2) parallel connectivity, in which the arguments of corresponding predicates also correspond. In addition, the selection of an alignment is guided by the *systematicity principle:* a system of relations connected by higher-

order constraining relations such as causal relations is preferred over one with an equal number of independent matches. Systematicity also guides analogical inference: people do not import random facts from base to target, but instead project inferences that complete the common system of relations (Bowdle and Gentner 1997; Clement and Gentner 1991).

Although analogy provides the strongest evidence for structure-mapping, alignment and mapping processes also apply in ordinary similarity (Gentner and Markman 1997; Markman and Gentner 1993; Medin, Goldstone, and Gentner 1993). For several years, we have run our main analogy simulation, SME (the Structure-mapping Engine), as an overall similarity engine. It forms matches at all levels, from object attributes to higher-order relations, and then sets about combining these into the best overall alignment. Because of the systematicity bias, relational alignments tend to win out, as in the case of examples (1) and (2). However, SME can also derive attributional solutions, as in (3), as well as interpretations that preserve both relational and attributional information, as in (4). (Indeed, SME's "favorite" class of comparisons, all else being equal, is literal similarity, in which there is a high degree of overlap in both object and relational information.)

Examples (l)-(4) all show structural consistency—one-to-one correspondences and parallel connectivity. Example (5) is more challenging. It belongs to a class of literary metaphors that lack clear one-to-one mappings and are characterized by many cross-weaving connections, with no clear way of deciding exactly which correspondences should hold (Centner 1982; Gentner, Falkenhainer, and Skorstad 1988). These kinds of metaphors seem to require processes such as metonymy and phonological matching in addition to alignment and mapping (see Fauconnier 1990, this volume; Fauconnier and Turner 1998 for further examples and analyses of complex metaphors.)

Structure-mapping makes a number of predictions that should follow if metaphors are processed like analogies. In the next two sections we consider the evidence. The first section asks whether extended metaphoric systems—which intuitively bear a strong resemblance to analogical systems—are in fact processed as analogical mappings between domains. The second section asks whether individual metaphors are processed by alignment and mapping.

Conceptual Metaphors as Extended Analogical Mappings

The presence of extended metaphoric systems is a striking feature of our language. People use analogies and metaphors from familiar concrete domains to discuss less familiar or abstract domains: e.g., flowing water for electricity (Gentner and Gentner 1983) or a long journey for marriage (Quinn 1987). Extended metaphors can evoke a whole system of knowledge, as when the computer metaphor, with its notions of encoding, storage, and retrieval, is applied to cognition.

Lakoff and his colleagues have documented many large-scale metaphoric systems in everyday language, such as the *argument as container* metaphor exemplified above, or the *love as a journey* metaphor: "The road was rough and steep but we carried on. . . . If we pull together we can surmount these hard times. We're having a rocky time and I'm not sure we're going to make it" (Lakoff and Johnson 1980; Lakoff and Turner 1989; Turner 1987). Many other domain-level metaphors have been described: marriage as a contract (Quinn 1987); the use of spatial terms to describe abstract dimensions such as economic prosperity or affective state (Nagy 1974); and the use of progeneration terms to express ideas like causality or preeminence, as in *mother of battles* (Turner 1987).

This research investigates the processing of extended metaphors during comprehension. In particular, we ask whether such metaphors are processed as mappings from one domain to the other during on-line comprehension—that is, whether participants establish and use a system-mapping between an initial domain (the base domain) and a second domain (the target domain) to process an extended metaphor. The presence of systems of metaphors between domains suggests that such metaphors are processed as systematic analogies (Gentner and Gentner 1983; Kittay 1987). But there are other possibilities. It could be that the metaphoric sentences in an extended metaphoric discourse are each processed separately and locally, with no connection across the sets of metaphoric phrases; or even that these seemingly metaphoric phrases are so conventionalized that their figurative meanings are directly available in the lexicon as alternate word senses. For example, the sentence *Ida gave Joe a great idea* could be processed as an on-line metaphor in which

give, which normally means to *cause a change in possession*, is mapped onto something like to *cause a change in cognitive state* (Gentner 1975; Gentner and France 1988). But it is also possible that the meaning to *cause a change in cognitive state* has become conventionalized as a possible meaning sense of *give,* in which case there need be no metaphorical mapping, merely a lexical look-up. Such cases may have originated as novel comparisons and become entrenched (the *cognitive archaeology* possibility). Thus, possibilities range from on-line generative mappings (the strongest claim) through the localist possibility that metaphors are simply processed and then discarded, to the weakest claim, that apparent metaphors might simply be processed as instances of multiple word senses (as polysemous word entries).

Localist Theories

There is little empirical evidence on the on-line processing of extended metaphors. With few exceptions, most theories of the processing of metaphor have emphasized local interactions between pairs of terms and ignored large-scale domain interactions. Many such theories are variants of the comparison view (Black 1962), in which finding the meaning of a metaphor involves finding the set of features that the base and target share (e.g., Malgady and Johnson 1980; Marschark, Katz, and Paivio 1983; Ortony 1979). For example, Ortony's (1979) influential salience imbalance theory of metaphor asserted that metaphoricity arises from an imbalance in the salience of the common features such that high-salient features in the base are matched with low-salient features of the target. In contrast, Glucksberg and Keysar's (1990) class-inclusion theory of metaphor, in explicit rejection of comparison theories, argues that metaphors are statements of category membership. A metaphor such as *my job is a jail* is understood by assigning the target of the metaphor *(my job)* to the category of which *jail* is the prototypical example: e.g., *confining institutions.* Glucksberg, McGlone, and Manfredi (1997) elaborated the category approach in their *attributive categorization* model. In this model, potential categories are generated and projected from the base while sets of modifiable dimensions are simultaneously identified in the target. For example, in the above metaphor, the base, *jail*, projects its *confining institutions* category while the target, *job*, yields its

modifiable dimensions: working conditions, degree of flexibility, and so on. The interpretation of the metaphor is thus an interaction between the category prototypically associated with the base and the dimensions that characterize the target. In this model, the base and target enter into processing in role-specific ways throughout the comprehension process.

The category approach can account for the finding that metaphors are asymmetrical: for example, the fact that *Some surgeons are butchers* and *Some butchers are surgeons* mean very different things, as noted by Ortony (1979) (see also Camac and Glucksberg 1984; Ortony et al. 1985). (However, other approaches can also account for asymmetry, as we will see.) The category-based approach is "localist:" it assumes that a metaphor conveys a categorical relation between a particular pair of terms. Thus this approach addresses single metaphors and not extended systems of metaphors.

Domain-Mapping Theories

One of the first theories aimed at large-scale domain-mappings was the domain-interaction hypothesis, proposed by Rumelhart and Abrahamson (1973) for analogies and by Tourangeau and Sternberg (1981, 1982) for metaphors (see also Tourangeau and Rips 1991). These theories used a multidimensional space model of mental representation and postulated that in metaphor, the dimensional structure of the base domain is mapped onto the dimensional structure of the target. Thus the metaphor *Brezhnev is a hawk* specifies an implicit mapping from the domain of birds to the domain of politicians, and states (with the appropriate dimensional substitutions) that Brezhnev's relative position among politicians—high in aggressiveness and political power—is the same as the hawk's among birds—high in ferocity and in strength. Tourangeau and Sternberg (1982) proposed that a good metaphor is one that, first, involves two very different domains and thus has high between-domain distance; and, second, shows low within-space distance between the base and target items in their very distant respective spaces: for example, *Brezhnev* and *hawk* occupy the same relative position in their domain spaces. Tourangeau and Sternberg found support for the theory's prediction that within-space closeness is positively correlated with aptness,

as well as a trend in favor of the second prediction—that between-domain distance should be positively correlated with aptness.

Kelly and Keil (1987) found evidence that metaphors can promote alignment across domains. Their participants rated concepts from two topic areas—say, periodicals and food—on semantic differential scales. Participants in the experimental condition then paraphrased and rated the aptness of four metaphors linking the two domains: for example, *The Wall Street Journal is the spinach of newspapers and magazines.* Then they again performed the same rating task. The experimental participants showed a shift in their ratings relative to a control group: pairs that made good metaphors became more similar in their semantic differential ratings and pairs that made poor metaphors became less similar. These findings are consistent with the claim that metaphors can induce large-scale domain-mappings.

The mental space representation of conceptual structures used in much of this earlier research has limited representational capabilities. It can represent dimensional structure, such as the ordering of members of a domain on attributes such as size or ferocity, but it cannot capture many other relations that are important in metaphor, such as causal structure, progeneration, and higher-order perceptual schemas (see Markman 1999). Nevertheless, these findings offer general support for the domain-mapping hypothesis.

A related approach from linguistics is Kittay and Lehrer's (1981) semantic field theory, which asserts that people understand metaphors through mapping the lexical fields that characterize the two domains. Metaphorically linking two domains alters one's view of one or both domains and this restructuring of domain(s) makes inferences about the target domain possible (Kittay 1987; Kittay and Lehrer 1981). For example, Kittay and Lehrer (1981) analyzed Wordsworth's poem comparing the history of Venice to the life history of a woman, in which the rise and decline of the city are likened to the youth and age of a noblewoman, Venice's inception in liberty (as a republic) is likened to a high-born birth, and so on.

Albritton, McKoon, and Gerrig (1995) found evidence suggesting that extended metaphors can result in schemalike long-term memory

as well as a trend in favor of the second prediction—that between-domain distance should be positively correlated with aptness.

Kelly and Keil (1987) found evidence that metaphors can promote alignment across domains. Their participants rated concepts from two topic areas—say, periodicals and food—on semantic differential scales. Participants in the experimental condition then paraphrased and rated the aptness of four metaphors linking the two domains: for example, *The Wall Street Journal is the spinach of newspapers and magazines.* Then they again performed the same rating task. The experimental participants showed a shift in their ratings relative to a control group: pairs that made good metaphors became more similar in their semantic differential ratings and pairs that made poor metaphors became less similar. These findings are consistent with the claim that metaphors can induce large-scale domain-mappings.

The mental space representation of conceptual structures used in much of this earlier research has limited representational capabilities. It can represent dimensional structure, such as the ordering of members of a domain on attributes such as size or ferocity, but it cannot capture many other relations that are important in metaphor, such as causal structure, progeneration, and higher-order perceptual schemas (see Markman 1999). Nevertheless, these findings offer general support for the domain-mapping hypothesis.

A related approach from linguistics is Kittay and Lehrer's (1981) semantic field theory, which asserts that people understand metaphors through mapping the lexical fields that characterize the two domains. Metaphorically linking two domains alters one's view of one or both domains and this restructuring of domain(s) makes inferences about the target domain possible (Kittay 1987; Kittay and Lehrer 1981). For example, Kittay and Lehrer (1981) analyzed Wordsworth's poem comparing the history of Venice to the life history of a woman, in which the rise and decline of the city are likened to the youth and age of a noble-woman, Venice's inception in liberty (as a republic) is likened to a high-born birth, and so on.

Albritton, McKoon, and Gerrig (1995) found evidence suggesting that extended metaphors can result in schemalike long-term memory

representations. Recognition of sentences from metaphorical passages was faster when both the target sentence and the preceding (priming) sentence shared a connection to the metaphor-based schema than when they did not. This research suggests that online mappings are created, but it does not test this directly, for participants made their recognition judgments *after* the reading metaphoric passages. Our research, presented below, takes up the issue of how extended metaphoric systems are processed on line. It uses reading time measures to determine whether participants are sensitive to the consistency of metaphoric mappings between domains during comprehension.

Lakoff and his colleagues argue for domain-level metaphors in a stronger sense than the foregoing researchers. They suggest that people possess large-scale conceptual metaphors such as "life as a journey" or "justice as balancing" (Lakoff and Johnson 1980; Lakoff and Turner 1989; Turner 1987, 1991). They have documented a large number of conceptual mappings that they propose are used to interpret metaphors and to inform their target domains. For example, Lakoff and Johnson (1980:90, 91) list examples of the *An argument is a journey* metaphor:

- We have *set out* to prove that bats are birds.
- *So far*, we've seen that no current theories will work.
- We will *proceed* in a *step-by-step* fashion.
- This observation *leads* to an elegant solution.
- He *strayed* from the main argument.
- Do you *follow* my argument?
- We have *covered a lot of ground* in our argument.

According to Lakoff (1990), metaphors are comprehended by invoking global mappings. Thus, this theory accounts for extended metaphor by postulating that people invoke a domain-level, prestored conceptual mapping when they encounter a local instance of the metaphor. To account for novel metaphors, Lakoff and Turner (1989) state that conceptual metaphors can be extended to new occurrences of existing domain-mappings (see also Gibbs 1990, 1994). Finally, a strong tenet of Lakoff's view is that metaphors do not draw on existing similarities, but rather *create* similarities by providing structure for the target domain.

Lakoff's claims about the psychology of metaphor have been controversial. For example, Murphy (1996, 1997) has noted difficulties in interpreting one of Lakoff's major theoretical claims, the *invariance hypothesis:* "Metaphorical mappings preserve the cognitive topology (that is, the image-schema) of the source domain, in a way consistent with the inherent structure of the target domain" (Lakoff 1993:215). One problem for the invariance hypothesis is the existence of multiple metaphors for the same target: e.g., *love is a journey, love is a disease, love is a fire.* Gibbs (1996) has defended the conceptual metaphor position on this last point, arguing that many abstract domains can be construed in several different ways, and therefore can accept metaphors from multiple base domains. However, a more serious challenge to the invariance principle is that the *same* base domain can provide different structures for different targets. For example, compare three metaphors that all use the base domain of *fire: knowledge is a fire* (one may pass the flame to others); *love is a fire* (its heat may consume the lover); *envy is a fire* (it burns upward toward its object, covering it with smoke). It is hard to argue that an invariant image-schema from *fire* is informing all these metaphors. It seems much more likely that our representations of *fire* include multiple schemas, and that which one is invoked depends on the alignment with the target domain. A related question is whether metaphors are understood by invoking preexisting conceptual metaphors, as Lakoff's theory suggests, or whether novel interpretations also occur. Finally, Murphy challenges Lakoff's claim that metaphors *create* meaning in the target, as opposed to reflecting parallels.

Despite these theoretical concerns, it is clear that Lakoff has identified an important phenomenon. Our research aims to capture the phenomena of large-scale mappings in a psychologically grounded account. We seek to explain how conceptual metaphors are processed and how they are learned. Are metaphors understood in terms of long-standing conceptual metaphors, or can mappings be constructed online, as most analogy theories assume? Structure-mapping theory suggests that metaphors are processed as structural alignments, based on some initial relational commonalties. Then further inferences are projected from the more concrete or familiar base to target. Thus, alignment highlights parallel structure (consistent with Murphy's position), and inference-

projection creates new knowledge in the target. This last fits with Lakoff's emphasis on new knowledge, but with the proviso that inference projection is guided by an initial alignment. This means that abstract domains are not structured de novo by concrete domains, but rather begin with some structure of their own and accept further structure from a commensurable concrete domain. Alignment serves to provide the needed constraint on possible conceptual metaphors.

Structure-mapping provides a natural mechanism for explaining how extended domain mappings are processed (Gentner 1982, 1983, 1988; Gentner and Clement 1988; Gentner and Markman 1997). Two key features that support extended mappings are the *systematicity bias* in interpretation and inference and the *incremental mapping* mechanism (Falkenhainer, Forbus, and Gentner 1989; Forbus, Gentner, and Law 1995; Gentner 1983; Gentner and Markman 1997). The systematicity bias—the preference for alignments that form deeply interconnected structures—fits with evidence that people naturally interpret analogy and metaphor by mapping connected systems of belief, rather than independent features. For example, Clement and Gentner (1991) asked people which of two common facts was most important to an analogy and, in another experiment, which new fact could be predicted from an analogy. For both tasks, participants were strongly influenced by systematicity: they were more likely to infer a new fact, and more prone to call a given fact important, if it was connected to a common causal structure. Systematicity is related to people's preference for relational interpretations of metaphors. Gentner and Clement (1988) found that adults interpreted metaphors such as *Plant stems are drinking straws* by invoking common relational systems (e.g., *They both convey liquids to nourish living things*) rather than object commonalties (e.g., *Both are long and thin*). Furthermore, adults (but not children) considered metaphors more apt when they found interpretations based on relational structure (Gentner 1988; Gentner, Rattermann, and Forbus 1993).

The second line of computational support for extended mappings is *incremental mapping.* An analogical mapping can be extended by adding further assertions from the base domain to the mapping (Burstein 1983; Novick and Holyoak 1991). For example, Keane and Brayshaw's (1988) [AM model simulates Keane's finding that subjects' initial mappings can

influence their subsequent correspondences. We have adapted this technique to create an incremental version of SME, called ISME. After creating a mapping between the initial base and target representations, ISME can extend the analogy by fetching further information from the base and seeking correspondences in the target that are consistent with the ongoing mapping. It thus enlarges the analogical mapping (Forbus, Ferguson, and Gentner 1994).

Assuming that a plausible process model can be provided for conceptual metaphors, there still remains the question of their psychological status. Are they cognitive mappings, or merely ways of talking? It is impossible to read the examples assembled by Lakoff and Johnson (1980), Turner (1987, 1991) Sweetser (1990), and others without feeling that at least some of them reflect psychologically real domain mappings. However, despite the intuitive appeal of the conceptual metaphor hypothesis, it would be rash to assume such a conclusion without more direct evidence. The perils of relying on intuition in interpreting metaphorical language are made clear in Keysar and Bly's (1995) study of the illusory transparency of idioms. They gave people archaic English idioms (for which the meanings are no longer current) in different contexts. People who heard *The goose hangs high* in the context of a sad story thought that it expressed sadness (a dead goose); those who heard it in a happy story thought it expressed happiness (a plentiful larder). More tellingly, both groups were confident that they would have arrived at the same interpretations without the story; they felt that their interpretation could be derived simply from the idiom. Keysar and Bly went on to suggest that the perceived transparency of an idiomatic expression increases with repeated use of an expression, and is not dependent on a genuine conceptual connection.

Researchers who study metaphor face the same difficulty as Keysar and Bly's subjects, of separating what they know to be the standard meaning of a conceptual metaphor from what the metaphor by itself suggests. These findings show that a feeling of transparency is not *by itself* evidence of a conceptual psychologically real mapping. A sense of transparency does not tell us whether the meaning is derivable from a systematic pattern or is simply learned by frequent lexical exposure. What is needed are techniques that go beyond speakers' intuitions.

In order to establish the conceptual role of domain mappings, we must first lay out the set of alternatives. There are at least four broad possibilities. The strongest possibility, as discussed above, is that metaphors *create* meaning (Lakoff 1990). In analogical processing terms, this would imply a purely *projective mapping:* the target domain is organized and structured in terms of conceptual systems borrowed from the more readily observable base domain. People actively use the metaphorical base domain to think about the target. The second possibility is *structural parallelism* in the domain representations, as suggested by Murphy (1996). Due to inherent similarities in the referent domains, parallel systems of semantic relations could evolve independently in two domains. Metaphors linking the two domains would then reflect structural alignment between the two parallel domain representations (Gentner and Markman 1997; Medin, Goldstone, and Gentnter 1993). In this case the two domains would share conceptual systems, but neither is derived from the other.

The third possibility is *cognitive archaeology:* systematic metaphors represent conceptual systems initially mapped from a particular base domain to a target domain, but which now exist as abstract systems that can apply to both domains. Such metaphoric relics would testify to the prior importance of a given analogical mapping in shaping the construal of the target domain in cultural history, they would not entail online mappings from the concrete domain during reasoning. However, they need not be purely local. To the extent that such systems preserve interconnections between their parts, they may still be processed as global systems in the target domain rather than in terms of individual assertions. The fourth and weakest possibility is *local lexical relations* (a kind of highly local cognitive archaeology). In this case, there are no large-scale systematic mappings; metaphors consist simply of individual polysemies and/or homophonies. For example, a term like *icy* could have concrete word senses, such as "made of frozen water; hard, cold," and also abstract word senses, such as "emotionally aloof, rigid, unyielding." A related possibility is that the two senses are stored as separate homophonic lexical entries. Either way the phenomenon would be purely lexical. It would entail neither large-scale structuring nor online mapping processes.

To settle these questions requires investigating the online processing of metaphors belonging to extended systems. Unfortunately, as noted above, most metaphor research has concentrated on individual metaphors—usually nominal metaphors of the form "An X is a Y." Much of this research has centered on testing the dual-stage hypothesis, proposed by Clark and Lucy (1975) and Kintsch (1974), which asserts that people first attempt to process linguistic material literally, and only if it cannot be understood literally do they try to process it nonliterally (Miller 1979; Searle 1976). Tests of this view typically compare processing of literal and figurative sentences (see Hoffman and Kemper 1987 for a review) and does not address the current question of whether and how people respond to systematic domain mappings.

Testing the Domain-Mapping Hypothesis

Gentner and Boronat set out to investigate whether extended metaphors are processed on-line as domain mappings (Boronat 1990; Gentner and Boronat 1992, in preparation; Gentner 1992). One potential obstacle to this kind of investigation is the fact that metaphorical language is often almost invisible. People use conventionalized metaphors such as *the weight of evidence* or *his spirits sank* without apparently noticing their metaphorical basis. Asking subjects for explicit judgments could interfere with natural processing. To get around this problem, we developed an indirect technique that does not require explicit choices. Our method is based on the phenomenon that mixed metaphors can be jarring and even humorous, as illustrated by these examples from the *New Yorker:*

This college is standing on a launching pad ready to flex its muscles

or

Reynaud was under the thumb of a mistress who was in the pocket of the pro-Axis party in France.

If such shifts of metaphor slow down processing, this suggests a disruption in the mapping process. Such a disruption would be consistent with the claim that people comprehend metaphors by setting up structurally consistent, systematic domain mappings. When two mappings are inconsistent, the resulting incongruity is then noticed. We used this mixed-metaphor technique to test the importance of consistency and

systematicity in the comprehension of extended metaphors. All the experiments followed the same logic, which can be illustrated with the followingexamples:

(1) Was Anna still boiling mad when you saw her?
No, she was doing a slow simmer.

(2) Was Anna still a raging beast when you saw her?
No, she was doing a slow simmer.

The initial sentences in (1) and (2) communicate roughly the same meaning. Both passages have the same final sentence. However, the last sentence in passage (2) involves a rather startling switch from the mapping set up in the first sentence (ANGER IS A BEAST) to the mapping in the final sentence (ANGER IS HEAT). In contrast, passage (1) uses the ANGER IS HEAT mapping throughout. Example (1), which maintains the same base to target mapping throughout, is *a consistently extended* metaphor. Example (2), which switches from one base domain to another, is *inconsistently extended:* the final sentence disrupts the metaphor set up by the first sentence. If participants take longer to read the last sentence in passage (2) than in passage (1), this will be support for the domain-mapping hypothesis, that people construct base-to-target mappings when they read extended metaphors, and that they naturally extend these mappings in a structurally consistent manner across connected base systems within the base and target.

We first describe two studies that used novel metaphors from existing conceptual mappings, as exemplified in box 6.1 (for more detail, see Boronat 1990; Gentner and Boronat 1992, in preparation; Gentner and Wolff 2000). The major contrast was between passages that used the same base throughout (*consistent*), and those in which the base was switched at a key point *(inconsistent).* The metaphor switch always occurred in the last sentence, for which reading times were collected. The consistent and inconsistent passages had the same story line and the same last sentence, but differed in the global metaphor used. For example, a short story about a debate was written using two different global metaphors[2] (see table 6.1). The consistent passage used the global metaphor A DEBATE IS A RACE: for example, . . . *he had to steer his course carefully in the competition.* The inconsistent passage used the

Box 6.1
Sample metaphoric passages utilizing novel metaphors (Gentner and Boronat 1992, in preparation)

Note: Metaphoric terms are in italics. They were presented in normal typeface to subjects. Passages were presented a sentence at a time.

Consistent: *A Debate is a Race*

Dan saw the big debate as a *race:* he was determined to win it. *He* knew that he had to *steer his course* carefully in the competition. His strategy was to go *cruising through* the initial points and then make his move. After months of debating practice, Dan knew how to present his conclusions. If he could only *keep up the pace,* he had a good chance of winning. Before long, he felt the audience was receptive to his arguments. Then, he *revved up* as he made his last key points. His skill left his opponent *far behind him* at the *finish line.*

Inconsistent: *A Debate is a War*

Dan saw the big debate as a *war:* he was determined to be victorious. He knew that he had to use every *weapon* at his command in the competition. He mapped out his strategy to insure he established a *dominant position.* After months of debating practice, Dan knew how to present his conclusions. If he could only *marshall his forces,* he had a good chance of winning. Before long, he felt the audience was receptive to his arguments. Then, he *intensified the bombardment* as he made his last key points. His skill left his opponent *far behind him* at the *finish line.*

Literal control:

Dan's topic in the big debate was "how to win a *race*": he had to be convincing. His first argument was on the proper way to *steer a course* in a competition. He argued strongly for *cruising through* initial laps and then making a move. After months of debating practice, Dan knew how to present his conclusions. If he could prove the need to *keep up the pace,* he had a good chance to win. Before long, he felt the audience was receptive to his arguments. His concluding remarks focused on *revving up* near the end of a race. His skill left his opponent *far behind him* at the *finish line.*

global metaphor, A DEBATE IS A WAR: for example, . . . *he had to use every weapon at his command in the competition.* For both passages, the last sentence used the RACE metaphor: for example, *His skill left his opponent far behind him at the finish line.* For the consistent passage, this represents a continuation of the global metaphor. However, for the inconsistent passage, the critical final sentence presents a switch between global metaphors: from A DEBATE IS A WAR to A DEBATE IS A RACE. The domain-mapping hypothesis predicts that the last sentence will be read more quickly following a consistent global metaphor than following an inconsistent global metaphor, because the former continues an established base to target mapping, whereas the latter disrupts it.

A literal control condition was also included to check for the possibility that such a mapping consistency effect, if observed, could be attributed to mere associative priming between the words in a passage and the words in the final sentence. The literal control passages contained all of the metaphoric terms of their matched consistent-metaphor passages, but these terms were used literally. Thus these subjects encountered the *terms* from the metaphoric base domain in the passage, but encountered the metaphor itself for the first time in the final test sentence. If these subjects read the final sentence as quickly as subjects in the consistent condition, it would suggest that any facilitation for the consistent condition over the inconsistent condition could be due merely to lexical priming.

The predictions are as follows. Localist metaphor theories, such as the class-inclusion theory of Glucksberg and Keysar (1990; Glucksberg, McGlone, and Manfredi 1997), would predict no difference between the two metaphoric conditions, since the key (metaphoric) sentence is the same. A modality-oriented view emphasizing differences between metaphoric and literal modes would also predict no difference between the two metaphor conditions, but would predict an advantage for both kinds of metaphor passages over the literal control passages. This is because the metaphoric last sentences should be faster to process after metaphoric than after literal material. Finally, the domain-mapping account predicts that test metaphors will be read faster in the consistent condition than in the inconsistent or literal control conditions.

Box 6.2
Sample metaphoric passage utilizing conventional metaphors

Consistent: *A Debate is a Journey*

Dan wanted to *guide* the audience *through* his debate speech. He did not *stray* from his *line* of argument. He showed that the opposition's arguments went off in the wrong *direction.* He won the debate because he *oriented* the judges to his interpretation.

Inconsistent: *A Debate is a War*

Dan wanted to *devastate* the audience with his debate arguments. He did not *desert* his *line* of argument. He *attacked* the opposition's arguments from the start. He won the debate because he *oriented* the judges to his interpretation.

Literal Control Passage:

Dan's *directions guided* him *through* the building to the debate room. He did not *stray* from the *lines* drawn on the map. He was well prepared to discuss problems with the opposition's arguments. He won the debate because he *oriented* the judges to his interpretation.

The results showed a mapping consistency effect: subjects read the last sentence significantly faster when it extended the existing mapping than when it drew on a new metaphoric mapping or followed a literal control passage. This finding supports the claim that processing extended metaphors involves alignment and mapping, and that people readily incrementally extend such mappings. These findings lend support to the claim that large-scale domain metaphors are psychologically real.

The evidence so far supports the domain-mapping hypothesis for novel metaphors.[3] But what about conventional metaphors? In two studies, Gentner and Boronat used passages that contained conventional metaphors, as exemplified in box 6.2. In many cases these metaphors came from the same global metaphors as the novel metaphors—e.g., DEBATE AS A RACE—but here the individual metaphors were

conventional. We reasoned that conventionalization might result in lexical storage—in metaphorical meanings coming to be stored with the base term. In this case, for highly conventional metaphors the metaphorical interpretation would simply be an alternate word sense. There should be no particular cost for switching global metaphors, since comprehension would simply involve finding the appropriate word sense. On this account, we would predict no advantage in time to read the last sentence for consistent metaphors over inconsistent metaphors.

The results were quite different from the first set of findings. Consistent with our speculations, when highly conventional metaphors were used, there was no apparent cost of shifting between global metaphors: subjects were not significantly slowed by a shift in the global metaphor. This supports the "alternate word sense" account for conventional metaphors. This finding is convergent with research suggesting that conventionalization results in a shift in metaphor processing from on-line active interpretation to retrieval of stored meanings (Bowdle and Gentner 1995, 1999). We return to this point later.

Summary

According to the domain-mapping hypothesis, people construct base-to-target mappings when they read extended metaphors, and extend these mappings in a structurally consistent manner across connected systems within the base and target. This predicts that metaphoric sentences will be read faster when they extend an ongoing mapping than when they require a new mapping, even when the conveyed meaning in the target is equivalent—the mapping consistency effect. In contrast, a localist account of metaphor—in which the passage metaphors are understood as local categorizations and then dropped—would predict no difference in reading time for the last lines of consistently and inconsistently mapped passages, since the same metaphors are being read and the meaning in the target is the same for both conditions.

We found evidence for a mapping consistency effect, but only for novel metaphors. Highly conventional metaphors were processed in a localist manner. We speculate that conventional metaphors behave like borderline idioms, with lexically stored figurative interpretations. This prediction is consistent with other evidence that highly familiar idiomatic

and metaphorical meanings are stored and processed at a lexical level (Blank 1988; Blasko and Connine 1993; Cacciari and Tabossi 1988; Gentner and Wolff 1997; Gibbs 1980, 1985, 1994; Gibbs and O'Brien 1990; Swinney and Cutler 1979).

The present results for novel metaphors are consistent with domain-mapping theories such as that of Rumelhart and Abrahamson (1973) and Tourangeau and Sternberg (1982) and with research suggesting that metaphors can be processed as large-scale conceptual systems (Allbritton, McKoon, and Gerrig 1995; Gibbs 1990, 1994; Nayak and Gibbs 1990; but see Glucksberg, Brown, and McGlone 1993). They argue against the kind of localist frame that is implicit in much current research. Theories that focus on single metaphors, such as Glucksberg and Keysar's (1990) class-inclusion theory of metaphor, cannot explain the links between extended metaphors, because they have no mechanisms for linking several discrete base-to-target mappings. Thus they cannot explain the pattern found for novel metaphors. However, as just noted, we believe a localist account may fit well with some kinds of conventional metaphors.

What Analogy Can Tell Us about the Processing of Individual Metaphors

Structure-mapping makes a number of predictions about the processing of individual metaphors that should follow if metaphors are processed like analogies. SME serves as a process model to motivate these predictions. The Structure-mapping Engine (SME) (Falkenhainer, Forbus, and Gentner 1989; Forbus, Ferguson, and Gentner 1994; Forbus, Gentner, and Law 1995) uses a local-to-global[4] alignment process to arrive at a structural alignment of two representations. Figure 6.1 shows SME's three stages of mapping. In the first stage, SME begins blind and local by matching all identical predicates in the two representations. Semantic similarity between predicates is captured through a decomposition into partial identities.[5] This initial mapping is typically inconsistent, containing many-to-one matches. In the second phase these local matches are coalesced into structurally consistent connected clusters (called *kernels*). The kernels are essentially partial mappings—connected sets of

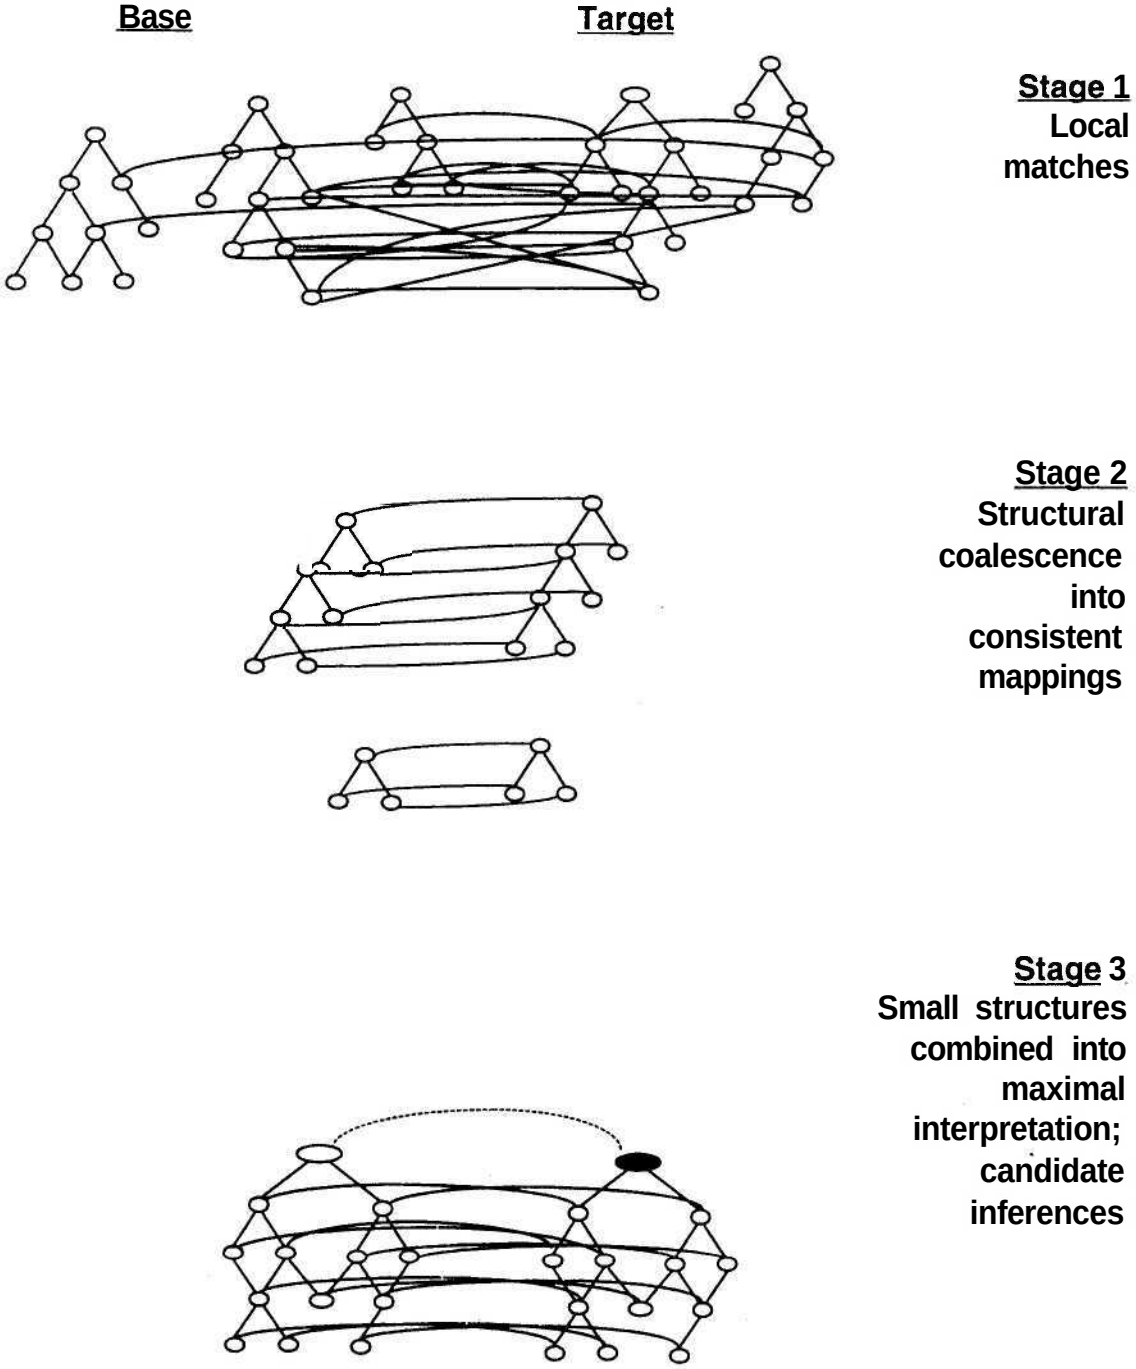

Figure 6.1
Overview of the algorithm used by the Structure-mapping Engine.

structurally consistent corresponding base-target pairs. They are given structural evaluations that depend not only on the sheer number of predicates but also on the depth of the system (Forbus and Gentner 1989).

Finally, in the third stage the kernels are merged into one or a few structurally consistent global interpretations (mappings displaying *one-to-one correspondences* and *parallel connectivity*). The challenge in finding the global interpretation is to find large, deep interpretations without having to exhaustively produce all possible interpretations, which would be psychologically implausible. SME uses a *greedy merge* algorithm (Forbus and Oblinger 1990) that operates in linear time.[6] It begins with the maximal kernel and then adds the largest kernel that is structurally consistent with the first one, continuing until no more kernels can be added without compromising consistency. It then carries out this process beginning with the second largest kernel to produce a second interpretation.

SME then produces a structural evaluation of the interpretation(s), using a kind of cascadelike algorithm in which evidence is passed down from predicates to their arguments. This method is used—both here and for the individual kernel evaluations mentioned earlier—because it favors deep systems over shallow systems, even given equal numbers of matches (Forbus and Gentner 1989). Up to this point, the processing has been a role-neutral process of alignment. Now, however, a directional inference process takes place. Predicates connected to the common structure in the base, but not initially present in the target, are projected as *candidate inferences* in the target. Thus, structural completion can lead to spontaneous unplanned inferences.

SME has several psychologically appealing features. First, it begins blindly, without needing to know the point of the comparison. Its alignment process begins by making large numbers of local matches, many of them mutually inconsistent. The global interpretations emerge by coalescing these matches in a manner that honors structural consistency and systematicity. Second, SME can derive two or three simultaneous interpretations of an analogy, capturing the finding that people can notice more than one interpretation of a metaphor. In particular, SME can derive literal and metaphorical interpretations simultaneously (We will return to this point.). Third, inference occurs as a natural outcome of

comparison, without special intention. This capacity to produce unanticipated inferences fits with the psychological intuition that inferences often arise unbidden in analogy, and may even surprise the reasoner.

This framework gives rise to several processing predictions. In particular,

- Metaphor comprehension begins with a symmetric (nondirectional) alignment process.
- If an alignment is found, then further inferences are directionally projected from base to target.
- Thus, directionality in metaphor comprehension arises *after* the initial stage of processing.

Directionality in Metaphor

According to the process model embodied in SME, metaphor comprehension begins with a symmetric alignment process: the representations of base and target are placed in correspondence and the largest and deepest consistent alignment(s) is found. At first glance, this claim of initially symmetric metaphor processing may seem far-fetched. After all, directional inference is one of the signature phenomena of metaphor: for example, in the metaphor *A rumor is a virus*, ideas such as contagion and sanitary habits are projected from the base concept, *virus*, to the target concept, *rumor*. Another symptom of metaphors' directionality is their nonreversability. For example, if the above metaphor is reversed, the result—*A virus is a rumor*—seems pointless. In other cases, reversing the terms radically changes the meaning. For example, *The acrobat is a hippopotamus* suggests a clumsy acrobat, while the reverse metaphor, *The hippopotamus is an acrobat*, suggests a graceful hippopotamus (Gentner and France 1988).

Ortony (1979) was the first to point out the importance of directionality for theories of metaphor. He proposed that metaphoricity arises through *salience imbalance*, when high-salient features of the base match with low-salient features of the target. Although empirical tests have not supported salience imbalance as the *cause* of metaphoricity (Gentner and Clement 1988; Tourangeau and Rips 1991), Ortony's insight that directionality is more important in metaphor than in literal similarity still stands. There is abundant evidence that reversing metaphors affects their

perceived aptness and interpretability (Gentner and Clement 1988; Glucksberg and Keysar 1990; Glucksberg, McGlone, and Manfredi 1997; Miller 1993; Ortony 1979; Ortony et al. 1985). For these reasons, most models of metaphor and analogy have assumed that processing is asymmetric throughout.

But in fact, there are at least three basic possibilities, as shown in figure 6.2. The strongest is that there is an initial *temporal asymmetry.* Processing begins with the base; after information is accessed or abstracted from the base, it is projected from the base to the target. A second possibility is that there is an *initial processing asymmetry.* Processing begins simultaneously with both terms, but is differentiated from the start in role-specific ways (Glucksberg, McGlone, and Manfredi 1997). The third possibility, predicted by structure-mapping and operationalized in SME, is *initial symmetry followed by processing asymmetry.* The initial stage is a role-neutral alignment stage; it is followed by a directional process of inference projection.

The first possibility, that processing temporally begins with the base, is explicitly or implicitly held by schema-projection models of analogy, such as Greiner's (1988) NLAG model or Keane and Brayshaw's (1988) IAM model. This processing order is also inherent in Glucksberg and Keysar's (1990) class-inclusion model of metaphor, in which metaphors are understood by finding or creating the category of which the base is the prototypical member and then applying this category to the target.

Gentner and Wolff (1997; Wolff and Gentner 1992) found evidence against temporal asymmetry using a priming method. We reasoned that if metaphor processing begins by accessing or deriving an abstraction from the base, and then projecting it to the target, then metaphor processing should be facilitated if the base term is presented just before the metaphor. In contrast, if the initial step is one of alignment, then there will be no special advantage for seeing the base over the target; either term will give a little advance information, and neither is sufficient to get very far.

To decide this issue, we measured subjects' time to interpret metaphors that were primed by either the base term or the target term. Initial projection models with temporal asymmetry predict faster comprehension given base priming than given target priming. The initial alignment

Three Processing Algorithms for Analogy, Similarity, and Metaphor

1. **Initial Projection: Temporal Asymmetry**

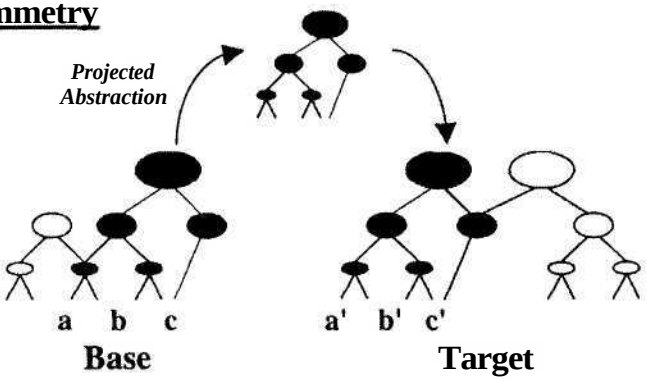

2. **Initial Abstraction; Processing Asymmetry**

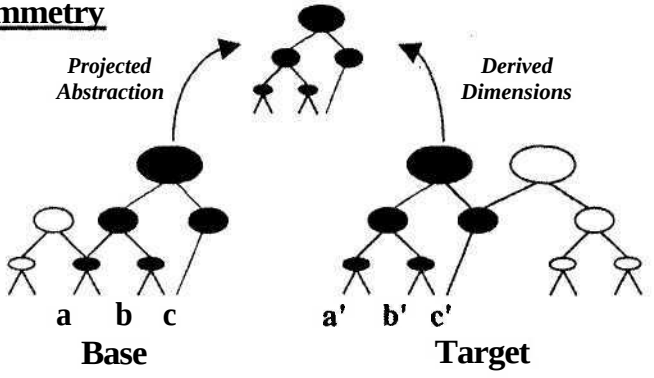

3. **Initial Alignment**

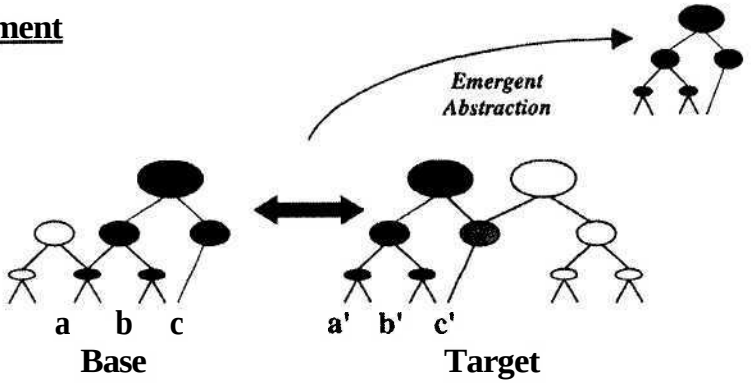

Figure 6.2
Three classes of process models. Model 1 shows initial projection of an abstraction from base to target. Model 2 shows initial projection of an abstraction from the base, along with dimensions from the target. Model 3 shows initial alignment, with subsequent extraction of the common abstraction.

model predicts no difference between base and target priming. Further, the initial alignment model predicts an advantage for high-similarity metaphors over low-similarity metaphors, because high (metaphoric or relational) similarity should facilitate initial alignment. The results were consistent with alignment-first processing. High-similarity metaphors were interpreted faster than low-similarity metaphors, and no base advantage was found across a series of experiments.[7]

Partly in response to these findings, Glucksberg, McGlone, and Manfredi (1997) proposed a more elaborated model (Model 2) in which there is a *processing* asymmetry, but not a temporal asymmetry, between the terms of a metaphor. In their *attributive category model*, processing begins simultaneously with both terms, but is differentiated from the outset in role-specific ways. The base term provides an abstract category that can be used to characterize the target, and the target provides dimensions along which it can be characterized. The base abstraction meets the dimensions derived from the target to form the interpretation. This proposal preserves Glucksberg and Keysar's position that metaphors are inherently class-inclusion statements, and that processing is directional from the start, but allows for Gentner and Wolff's finding of no general base advantage in priming.

For example, in *My surgeon is a butcher*, a category such as *one who cuts flesh crudely* is derived from the base term, *butcher*, while a set of dimensions, such as *skill level* and *income level*, is derived from the target term, *surgeon*. The base category is combined with the target dimensions to produce the metaphor interpretation. As evidence for role-specific processing, Glucksberg and his colleagues showed that metaphors were faster to be comprehended when primed by *unambiguous* bases—that is, bases that uniquely exemplify a particular category—than by *ambiguous* bases, and by *high-constraint* targets—that is, targets that strongly constrain the attributive dimensions along which they can be characterized—than by *low-constraint* targets. However, this study did not demonstrate that these effects were role-specific. That is, there was no test of the equally important *negative* prediction—that low ambiguity in the target and high constraint in the base would *fail* to facilitate comprehension. Lacking such a demonstration, it remains possible that degree of ambiguity and degree of constraint are both simply aspects of the *informa-*

tivity of terms with respect to the metaphoric interpretation; and that high-informative terms (not surprisingly) facilitate comprehension of metaphor to a greater degree than low-informative terms.

At this point, it seemed that the two groups had reached a kind of priming stalemate. Recall that in structure-mapping, the initial stage of metaphor processing is a symmetric alignment process; directional inference projection does not occur until after the representations are aligned (Bowdle and Gentner 1997; Falkenhainer, Forbus, and Gentner 1989; Forbus, Gentner, and Law 1995; Gentner and Markman 1997; Gentner and Wolff 1997; Wolff and Gentner 1992, 2000, in preparation). For example, given the metaphor *My surgeon is a butcher,* an initial (symmetric) alignment process would yield the common system *one who cuts flesh.* Then role-specific inference processes would project further ideas connected to this schema from the base to the target: for example, the idea of *cutting crudely and without regard for the health of the flesh.* Thus the question was how to test between processing that begins role-neutral and ends role-specific (as in structure-mapping) and processing that begins with two simultaneous role-specific processes (as in the attributive category model), that is, between Models 2 and 3 in figure 6.2.

What was needed was a better way to probe initial stages of processing. Wolff and Gentner (2000) found such a technique by adapting a metaphor interference technique originally pioneered by Glucksberg, Gildea, and Bookin (1982) and extended by Keysar (1989). In Glucksberg et al.'s task, participants were simply asked to judge statements as literally true or false. There were three basic kinds of statements: true class-inclusion statements (e.g., *Some birds are robins*), false class-inclusion statements (e.g., *Some birds are apples*), and metaphorical statements (e.g., *Some jobs are jails*). Subjects were told to respond "true" only for the first class, the literally true statements. Response times were recorded. As expected, participants could speedily classify correct class-inclusion statements as true and false statements as false. The key finding was that participants found it hard to reject metaphors: they were slower to respond "false" to metaphors than to ordinary false statements. The fact that metaphoric meaning interfered with literal true-false judgments was an important finding in the history of metaphor, for

it implies that processing of metaphorical meanings begins *before* a true-false literal judgement has occurred.

For present purposes, the beauty of Glucksberg et al.'s interference effect is that it appears to tap into the initial processing stages. Metaphoric interference effects appear at about 1,200 milliseconds, well below the two to four seconds typically required for full comprehension of metaphor. Thus the metaphor interference technique offers a way to probe very early metaphoric processing.

Our technique was simply to repeat the task used by Glucksberg, Gildea, and Bookin (1982) and Keysar (1989), with one key alteration: We included reversed metaphors—e.g., *Some handcuffs are contracts*—as well as forward metaphors—e.g., *Some contracts are handcuffs.* (See Wolff and Gentner, 2000, for details.) According to Glucksberg, McGlone, and Manfredi's (1997) attributive category account, initial processing is role-specific. If the terms are in reversed order, the sentence should simply seem anomalous, because the category provided by the base will not fit with the dimensions provided by the target. (All the metaphors were pre-tested to ensure that they were strongly directional, so that only the forward direction made sense.) Thus, metaphoric interference is expected only when the terms are in the forward order; a reversed ordering of the terms should lead to no more interference than an ordinary false statement. In contrast, if the initial process is structural alignment, then the early stages should be role-neutral; only later should subjects detect the bizarreness of the reversed metaphors. Hence, structure-mapping predicts that metaphoric interference (1) should be independent of the order of the terms; and (2) should increase with the similarity between base and target (because, as noted above, high similarity facilitates alignment).

The results were fairly dramatic. In the first study, we replicated Glucksberg et al.'s interference effect: Metaphors took longer to reject than ordinary false statements, indicating early processing of metaphoric meaning. The key finding, however, was that reversed metaphors showed just as much interference as forward metaphors. (Time to say *false* was 1,118 msec for forward and 1,111 msec for reversed metaphors; these times did not significantly differ from each other, but both were longer

than the 1,064msec found for ordinary false statements.) That interference effects were independent of the order of the base and target terms is strong evidence for initial role-neutral processing, even for highly directional metaphors.

In the second study, we again probed initial processing, this time using metaphors that seemed particularly apt for revealing role-specific processing—namely, high-conventional metaphors, whose base terms have stock metaphorical senses associated with them. We reasoned that such metaphors should be particularly likely to show directional role effects. As before, initial projection theories predict strong effects of direction, and no early effects of similarity. Initial alignment theories such as structure-mapping predict strong early effects of similarity and not of direction.

The results were again striking. We again found symmetric interference effects. Even highly conventional metaphors showed no direction dependence in interference effects.[8] Further, interference effects occurred only for high-similarity metaphors. This result is exactly what would be expected if the initial processing of a metaphor were structural alignment. These results are evidence for initial alignment theories and against initial projection theories.

In the third study, we verified that the metaphors were directional when processed to completion. According to the structure-mapping model, the initial alignment process is followed by directional projection of inferences. Thus metaphoric directionality should emerge if people are allowed to complete the comprehension process. To test this prediction, we gave subjects the same metaphors as in the prior studies and asked them to judge whether or not the metaphor was comprehensible. If later processing is directional, as predicted by structure-mapping theory, then forward metaphors should be rated as more comprehensible than reversed metaphors.

This result is exactly what was found. Participants judged forward metaphors (M = .75) to be comprehensible about twice as often as reversed metaphors (M = .37). Further, forward metaphors were comprehended in less time than reversed metaphors (M = 1,644ms for forward, M = 1,778 ms for reversed). The third prediction of the structure-mapping model, that high-similarity metaphors should be easier to comprehend

than low-similarity metaphors was also borne out. High-similarity metaphors were more likely to be judged comprehensible than low-similaritymetaphors.

In sum, the results support a process model in which an early symmetrical alignment process is followed by later directional processing. Early processing of metaphors, as tapped by the interference effect, is symmetrical. However, when full processing is allowed, a pronounced asymmetry appears between forward and reversed metaphors. Overall, the pattern fits the structure-mapping claim of initially symmetric processing followed by later directional projection of inferences.

Implications for Models of Analogy

Models of analogical processing, like models of metaphor, differ in whether they begin with a directed projection process, followed by matching and verification, or with a symmetric matching process, followed by directed projection. The former class includes explanation-based-learning models of analogy (Kedar-Cabelli 1988), abstraction-based models such as Greiner's (1988) NLAG, Keane and Brayshaw's (1988) IAM, and possibly Hummel and Holyoak's (1997) LISA model, which operates in driver-recipient mode rather than by alignment,[9] as well as incremental projection models such as that of Keane, Ledgeway, and Duff (1994). The latter class includes alignment-first models: SME (Falkenhainer, Forbus, and Gentner 1989) and ACME (Holyoak and Thagard 1989). The findings reviewed here are strong evidence for initially symmetric, role-neutral processing and, more generally, for alignment-based models of analogy and metaphor.

The Career of Metaphor

The alert reader may have noticed that conventional metaphors have differed from novel metaphors in several ways in the studies described so far. In the first section, novel metaphors, but not conventional metaphors, showed mapping consistency effects in processing extended metaphors (Gentner and Boronat 1992). In the second section, the one exception to our general finding of equivalence between base and target priming was that highly conventional, low-similarity metaphors showed

a base priming advantage (Gentner and Wolff 1997) (See note 7). In this section we lay out a theory of how conventional metaphoric meanings arise and how the representation and processing change as a metaphor progresses from novel to conventional.

We propose an account of metaphor and polysemy that we call the *career of metaphor* hypothesis: Novel metaphors are processed as structural alignments between the concrete or literal representations of the base and target, but as repeated comparisons are made, the metaphorical meaning is gradually abstracted and comes to be associated with the base term (Bowdle and Gentner 1995, 1999, in preparation; Gentner and Wolff 1997, 2000; Wolff and Gentner 1992, 2000, in preparation). We know from research on analogical problem solving that the alignment of two relationally similar situations can lead to the induction of domain-general problem schemas (e.g., Gick and Holyoak 1983; Loewenstein, Thompson, and Gentner, 2000; Novick and Holyoak 1991; Ross and Kennedy 1990). We believe that similar forces are at work during metaphor comprehension. The process of structural alignment allows for the induction of metaphoric categories, which may come to be lexicalized as secondary senses of metaphor base terms (Bowdle and Gentner 1995, 1999, in preparation; Gentner and Wolff 1997).

This kind of progressive abstraction can be computationally modeled by storing the common schema that SME derives from a comparison. We have used an extension of SME called SEQL to carry out this kind of schema abstraction (Skorstad, Gentner, and Medin 1988). In this model, the common schema that SME derives from carrying out a comparison is stored as an abstraction and carried forward. It can then be aligned with further exemplars. This process typically results in a progressively more abstract relational structure, with fewer and fewer surface details. Such abstractions could serve as metaphoric category representations. We suggest that the deriving and retaining of structural abstractions is the basic mechanism by which metaphors become conventionalized.

When a metaphor is first encountered, both the target and base terms refer to specific concepts from different semantic domains, and the metaphor is interpreted by aligning the two representations and importing further predicates from the base to the target. One result of this

mapping is that the common relational structure that forms the metaphor interpretation will increase in salience relative to nonaligned aspects of the representations. If the same base term is repeatedly aligned with different targets so as to yield the same basic interpretation, then the highlighted system may become conventionally associated with the base as an abstract metaphoric category.

The gist of these claims is illustrated in figure 6.3, which shows the evolution from novel to conventional metaphor. *Novel metaphors* involve base terms that refer to a domain-specific concept, but are not (yet) associated with a domain-general category. They are interpreted as comparisons: direct structural alignments between the literal base and target concepts. *Conventional metaphors* involve base terms that refer both to a literal concept and to an associated metaphoric category. At this point in its evolution, the base term is polysemous, having both a domain-specific meaning and a related domain-general meaning. For

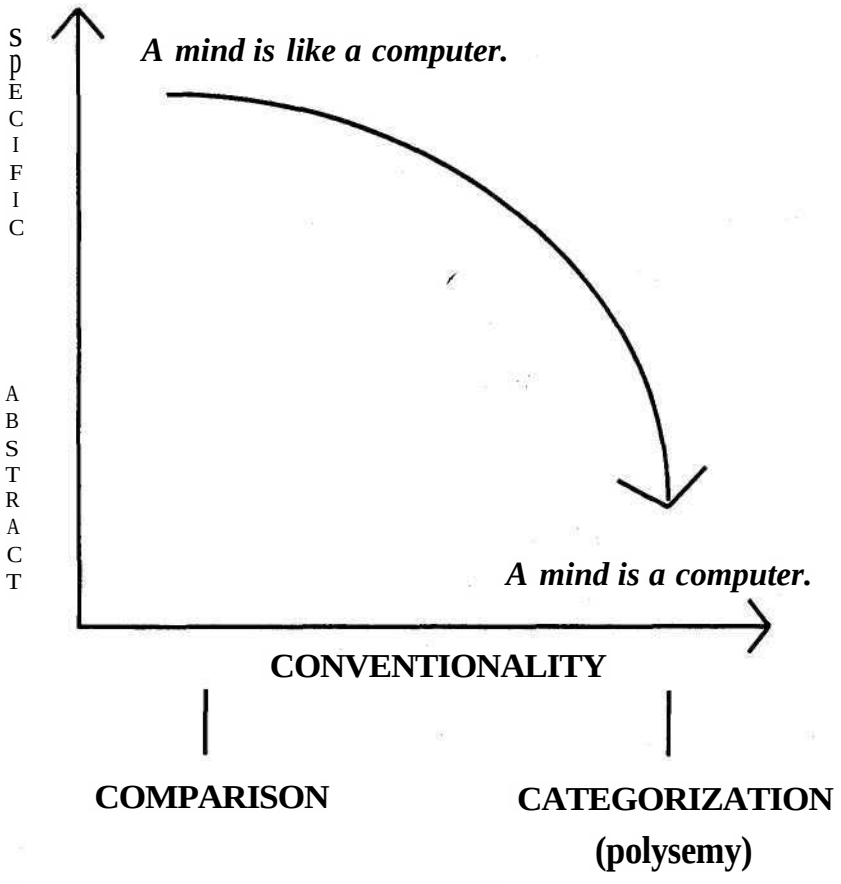

Figure 6.3.
The career of metaphor from novel to conventional metaphor.

instance, Sweetser (1990) suggests that many standard terms for mental processes are derived from terms for sensory perception: for example, *discernment*, *focus*, *point of view*, and so on (see also Holyoak and Thagard 1995; Murphy 1996).

Implications for Metaphor Comprehension

Research on metaphor comprehension has usually treated metaphor as an undifferentiated class. However, a number of theorists have recently argued that how a metaphor is processed may depend on its level of conventionality (e.g., Blank 1988; Blasko and Connine 1993; Giora 1997; Turner and Katz 1997). Our account of the relationship between metaphor and polysemy is in line with these claims. Specifically, we suggest that novel metaphors are processed by aligning the literal senses of both terms. In contrast, conventional base terms are often polysemous, with both literal and metaphoric meanings. Thus they may be interpreted either by aligning the target concept with the literal base concept, or by aligning the target concept with the base's abstract metaphoric category. We could think of the first as comparison and the second as categorization.[10]

Thus, the career of metaphor hypothesis predicts that as metaphors become increasingly conventional, there is a shift from comparison to categorization (Bowdle and Gentner 1995, 1999, in preparation; Gentner and Wolff 1997, 2000; Wolff and Gentner, forthcoming). This is consistent with recent proposals that the interpretation of novel metaphors involves sense creation, but the interpretation of conventional metaphors involves sense retrieval (e.g., Blank 1988; Blasko and Connine 1993; Giora 1997; Turner and Katz 1997); and likewise for idioms (Cacciari and Tabossi 1988; Gibbs 1980, 1994; Swinney and Cutler 1979; Williams 1992). On the present view, the senses retrieved during conventional metaphor comprehension are abstract metaphoric categories.

Bowdle and Gentner (1995, 1999, in preparation) tested five predictions of the career of metaphor account. The first two predictions took advantage of a proposed form-process link whereby the simile form (e.g., *A child is like a snowflake*)invites comparison and the metaphor form (e.g., *A child is a snowflake)* invites categorization (Bowdle's grammatical concordance assumption). This predicts that giving subjects the same

pairs of terms in either simile or metaphor form should invite different forms of processing. Specifically, the predictions are: (1) novel figurative statements should be easier to process as similes than as metaphors. This is because the simile form invites comparison processing whereas the metaphor form invites subjects to search for a prestored category (which in the case of novel statements does not exist). (2) novel figurative statements should be preferred in the simile (comparison) form rather than the categorization form. Both these predictions were borne out.

To test the prediction that novel figuratives should be easier to process in the comparison form (e.g., *A mind is like a kitchen)* than in the categorization form *(A mind is a kitchen)* Bowdle and Gentner gave subjects a timed comprehension task. We found that novel figuratives were read faster in comparison (simile) form than in categorization (metaphor) form. The reverse was true for conventional figuratives; these were processed faster in categorization form *(An opportunity is a doorway)* than in comparison form *(An opportunity is like a doorway).* To test the second prediction, that novel figuratives should be preferred in the comparison form, we asked subjects to rate their preference for simile vs. metaphor form for a range of figurative statements. As predicted, participants strongly preferred the simile form for novel figuratives; they showed no preference between the simile and metaphor forms for conventional figuratives.

We also found evidence for two further predictions, one rather humdrum and one quite surprising, at least from the point of view of class-inclusion theories. The humdrum result is that people rated novel figurative statements as more metaphorical than conventional figuratives. This follows from the idea that the perception of metaphoricity reflects active construction of new alignments and inferences, such as occur for novel metaphors. The more surprising result is that people considered similes *more metaphorical* than metaphors overall. Specifically, similes were rated more metaphorical than metaphors for conventional figuratives, and equally metaphorical for novel figuratives. We believe this is because similes invite a fresh alignment between literal senses, even for conventional metaphors. For example, the stock metaphor *Fred is a gorilla* calls forth the standard "large, aggressive, and dangerous" abstraction; but in simile form—*Fred is like a gorilla*—it invites ideas

such as "gentle despite strength, limited to a specific environmental niche, and vulnerable to extinction."

Finally, we directly tested the process account postulated in the career of metaphor hypothesis. We asked whether we could achieve conventionalization in vitro, by giving participants several parallel novel figurative statements using the same base term (Bowdle and Gentner 1999; in preparation). Participants were given two novel similes and asked to create a third that was similar in meaning to the first two: for example,

Doubt is like a tumor.
A grudge is like a tumor.
———is like a tumor.

Each participant saw eight such sets, as well as other sets containing literal similarity statements, metaphors, and categorizations. Then participants were given test sentences containing the same bases in both simile and metaphor form and asked to say which they preferred—for example,

An obsession is a tumor. *An obsession is like a tumor.*

Subjects who received the in vitro conventionalization condition (i.e., those who saw two novel similes) were significantly more likely to prefer the metaphor form than subjects who simply saw the test sentences with no prior manipulation, or who had seen the same base terms in literal similarity statements. It appears that carrying out a series of parallel figurative alignments led to the creation of an abstraction that favored the metaphor form.

Metaphor and Polysemy

Many linguists have proposed a link between metaphor and polysemy. A common proposal is that lexical extensions (e.g., the use of *see* in 7 *see your point*) are comprehended via a mapping from a (typically concrete) domain of experience to another (typically abstract) domain of experience (e.g., Lakoff and Johnson 1980; Lehrer 1990; Sweetser 1990); the concrete domain is invoked to explain phenomena in the abstract or less familiar domain. On this view, the metaphoric meaning of a polysemous word is understood directly in terms of its literal meaning.

Bowdle and Gentner (1995, 1999, in preparation) suggested a related but somewhat different relationship between metaphor and polysemy,

one that follows naturally from considerations of analogical processing. The process of structural alignment allows for the induction of metaphoric categories, which then may be lexicalized as secondary senses of metaphor base terms (Bowdle and Gentner 1995, 1999, in preparation; Gentner and Wolff 1997). Such processes provide ways of creatively extending the lexicon (Gerrig and Gibbs 1988; Lehrer 1990). On this account, it is not necessary to return to the base to process conventional metaphors. Indeed, the metaphoric sense may persist even after the literal sense disappears, as discussed below.

From Comparison to Categorization

The career of metaphor hypothesis can be related to the proposal that metaphor is a species of categorization (e.g., Glucksberg and Keysar 1990; Glucksberg, McGlone, and Manfredi 1997; Honeck, Kibler, and Firment 1987; Kennedy 1990). As discussed earlier, Glucksberg and his colleagues have proposed that metaphors are processed as class inclusions: the base concept is used to access or derive an abstract metaphoric category to which the target concept is assigned. Our results suggest that this account is fundamentally wrong for novel metaphors, but that it may apply well to some conventional metaphors. Highly conventional metaphor bases are just those that have associated metaphorical abstractions—*dual representations* in Glucksberg and Keysar's (1990) terms.

What about intermediate stages of conventionalization? One appealing account is a race model (see Cacciari and Tabossi 1988 for idioms; Gentner and Wolff 1997 for metaphor). In this model, the literal and figurative interpretations are accessed in parallel and the first one to achieve a satisfactory interpretation wins. If it is faster to retrieve and apply a stored figurative interpretation than to create a new interpretation by aligning a base and target representation, then the stored figurative interpretation will prevail. The more conventionalized the base term, the more likely it is that the stock interpretation will prevail. The more similar the literal representations of the base and target, the more likely it is that the process will be one of alignment of literal senses.

The centrality of comparison. What all this implies is that comparison is the fundamental process that drives metaphor. Novel metaphors are

understood *only* by comparison. Conventional metaphors can be understood by accessing stored abstractions, but these metaphoric abstractions are a *product* of past comparisons. Thus, although conventional metaphors can be seen as category-inclusion statements, the categories themselves are derived from prior alignment processes. To use one of Turner's procreative metaphors, one could say that comparison begets categorization.

There are three further reasons to emphasize comparison and alignment as central in metaphor interpretation. First, Bowdle and Gentner's results suggest that conventional metaphors can readily be processed either in comparison (simile) form or in categorization (metaphor) form, whereas novel metaphors are far more understandable in comparison form. Comparison is thus the more universal process for metaphor. Second, the career of metaphor account suggests that the class-inclusion stage occurs fairly late in the life of a metaphor. Third, even for conventional metaphors the *process* of comprehension may be structural alignment. The results of the Wolff and Gentner (2000) metaphor interference studies suggest that even for conventional metaphors, the initial process is alignment: even for highly conventional metaphors, reversed and forward metaphors were indistinguishable in initial processing. Thus, the career of metaphor from comparison to categorization may involve a change in *representation* rather than a change in *process*.[11]

Metaphor Senescence and Death: Conventionality and Idiomaticity

Research into real-time idiom processing can shed light on the processes used in highly conventional "stock" metaphors. One prominent view of idioms is that their figurative interpretations are stored in memory as single lexical items (Cacciari and Tabossi 1988; Ortony et al. 1978). This *idiom as dead metaphor* hypothesis posits that, at least in the lexicon, idioms' stored figurative meanings are interpretations stripped of their base to target mappings. Accessing the figurative meaning of the idiom *pop the question,* then, would tell one that the idiom referred to a marriage proposal; it would not tell one about the action of popping or the nature of questions.

Gibbs, Nayak, and Cutting (1989) distinguish between *decomposable idioms,* whose meanings can be deduced from their parts (e.g., *pop the*

question), and *nondecomposable idioms*, whose meanings cannot be so deduced (e.g., *chew the fat*). Nondecomposable idioms seem relatively identifiable as idioms, while decomposable idioms, which may be linked to conceptual metaphors, seem less so. Gibbs, Nayak, and Cutting's findings suggest that people do not always analyze the literal meanings of idioms; rather, they perform a componential analysis that assigns a figurative interpretation to each of an idiom's components. Thus, someone reading *pop the question* would ascribe to it the meanings *suddenly say* and *a marriage proposal.* No explicitly "literal" interpretations of the idiom need be produced. Likewise, McGlone (forthcoming) suggests that idioms may be psychologically processed in units smaller than the whole phrase.

However, Gibbs and O'Brien (1990) have contested the claim that idioms are entirely stripped of their base-to-base target mappings. In a study analyzing protocols, they found considerable regularity in participants' descriptions of their mental images of idioms (e.g., *lose your cool*), a regularity lacking in other participants' descriptions of similar literal phrases (e.g., *lose your wallet*), or in descriptions of figurative interpretations (e.g., *to get angry).* They interpreted this regularity as indicating that these idioms remain connected to conceptual metaphors such as ANGER IS HEAT or THE MIND IS A CONTAINER (Lakoff and Johnson 1980). While their results suggest that idioms *can* regain their base-to-target mappings, Gibbs and O'Brien explicitly reject the idea that conceptual metaphors play an on-line role in normal idiom processing, given (a) that idioms are understood so rapidly and (b) the evidence reviewed earlier that the figurative meanings of idioms are stored in memory (Cacciari and Tabossi 1988).

In summary, the *career of metaphor* proposal traces metaphors from early alignment and mapping between literal meanings (the novel metaphor stage) to a later stage of dual representation in which the metaphor may be understood either by a novel alignment or by accessing an abstract representation (the conventional metaphor stage), to a stage in which the metaphoric representation has become a standard word sense for the base (the polysemy stage). At this point, the sense of metaphoricity disappears, and only polysemy remains. Sometimes, a still further stage occurs, in which the literal meaning disappears entirely,

leaving only the (formerly) metaphoric sense as the meaning of the base term.

General Discussion

Analogy research offers a set of psychologically and computationally tested processes that can serve to illuminate both the processing of individual metaphors and extended systems of metaphors. Further, structure-mapping offers a way in which metaphorical and literal comparisons can be captured within a single mechanism.

Addressing the Classic Problems of Metaphor

We have suggested that structure-mapping can serve as an account of metaphor processing. To make good on this argument we must ask whether this framework can handle the central phenomena of metaphor, and how it deals with the classic challenges. The following four issues are points on which any theory of metaphor must take a stand:

1. *Metaphor recognition:* how and when metaphoric (as opposed to literal) processing is initiated
2. *Metaphoric meaning:* do metaphors create meaning or do they simply reflect parallels
3. *Metaphoric induction:* how (and how *much)* property inference occurs
4. *Directionality:* why metaphors are preferred in one direction over the other

The approach from analogy to metaphor offers some solutions.

1. *Metaphor recognition: How and when is metaphoric processing initiated?* A thorny problem in metaphor has been how people distinguish metaphors from literal statements so that they know to initiate metaphoric processing. What signals us to initiate metaphoric processing for *Genghis Khan was a butcher*, but not for *My neighbor is a butcher*? The classic proposal that metaphor is initiated only after literal processing fails has been rejected, but no new consensus has emerged. In SME, this problem does not arise. As noted above, SME normally produces at least two simultaneous interpretations of a comparison. In par-

ticular, it can derive literal and metaphorical interpretations simultaneously. If I tell you *My neighbor is a butcher,* for example, you might entertain both the possibility that his profession is slaughtering animals and the metaphoric possibility that he butchers something else—perhaps hedges? Which interpretation you choose will depend on (a) which has the largest and deepest common structure, as discussed above; (b) which best fits other knowledge about the target; and (c) which is more relevant to the current context. Likewise, if given *Genghis Khan was a butcher,* you might briefly imagine a literal interpretation—that his profession was slaughtering animals—but prior knowledge would quickly lead you to reject it in favor of the correct metaphorical alignment—that he slaughtered human beings.

2. *Metaphoric structure: Do metaphors create meaning, or merely reflect structural parallels?* The account offered by Lakoff and his colleagues (Lakoff 1990; Lakoff and Johnson 1980; Lakoff and Turner 1989) has it that metaphors create meaning: conceptual systems are projected from one domain to another, typically more abstract, domain. In contrast, Murphy (1996) proposes that metaphors do not provide new structure for the target domain, but rather reflect structural parallelism between two domains. Structure-mapping incorporates aspects of both these accounts. It begins with structural alignment (as in Murphy's account) but also allows for further candidate inferences—that is, for a constrained kind of meaning creation.

3. *Property induction. If metaphor is a comparison, then how does one account for property induction?* The fact that metaphors can convey new information is sometimes used to argue against comparison theories, on the grounds that merely highlighting common properties does not allow for new information (Glucksberg and Keysar 1990; Glucksberg, McGlone, and Manfredi 1997). But, as discussed above, inference projection naturally arises out of alignment. Indeed, it is far from clear that a pure projection theory can be made computationally tractable.

4. *Directional asymmetry: How can a comparison approach account for the strong directionality of metaphors?* People show strong directional preferences in metaphor. Thus, people tend to think that (1) is an intelligible metaphor, but that (2) is not:

(1) Some jobs are jails.

(2) Some jails are jobs.

As discussed earlier, the strong directionality of metaphors is a classic challenge to comparison models, and has been used to argue that metaphors are essentially class-inclusion statements rather than comparisons. Our evidence indicates that asymmetry in comprehension arises after the initial alignment stage. Inferences are directionally projected from the base to the target. This predicts that both speaker and hearer should prefer to have the more informative term in the base position.

Bowdle and Gentner (1997) verified that informativity can determine directionality, even for literal comparisons, which are less asymmetric than metaphors. Subjects read two brief passages, which were similar except that one passage contained a systematic causal structure linking the events and the other did not. When asked to generate any inferences they chose between the passages, subjects were overwhelmingly more likely to project inferences from the more systematic passage to the less systematic one. In another study, subjects were given two comparison statements—"A is similar to B" or "B is similar to A"—and asked to choose which direction they preferred. They consistently preferred the more systematic passage as base, and considered this direction more informative.[12] These findings establish a connection between inferential potential and preferred direction, as predicted by structure-mapping.

We suggest that systematicity imbalance can explain the directional asymmetry of metaphor. This would fit with the preponderance of embodiment metaphors noted by Lakoff and his colleagues. We rely heavily on mappings from experiential domains such as spatial relations and bodily force dynamics, because our models in these domains are sufficiently systematic to provide inferential structure for other domains.

Implications for Metaphor Research

Metaphor is related on the conventional side to idiom and on the novel side to analogy. But metaphor research has focused disproportionately on conventional metaphors. Indeed, Bowdle and Gentner (in preparation) surveyed the metaphors used in a sample of psychology research papers and found that most were of high conventionality. The current

findings underscore the importance of conventionality as a factor in experiments. We need to reexamine claims about metaphoric processing that are based only on conventional metaphors.

Implications for Models of Analogy

This research has both local and global implications for models of analogical processing. At the local level, Wolff and Gentner's findings offer support for alignment-first models such as SME over projection-first models such as IAM. We found evidence for initial role-neutral processing. Even for strongly directional metaphors, forward and reversed orders appear to be initially equivalent in processing, consistent with SME's initial symmetric-alignment process. Our finding of directional superiority at longer deadlines is consistent with SME's later process of directed projection of inferences. In contrast, models that begin by directionally projecting information from base to target, such as IAM (Keane and Brayshaw 1988) and LISA (Hummel and Holyoak 1997) (though see note 9), should predict directional superiority from the start. Thus the initial equivalence of forward and reversed metaphors is a serious challenge to these models. Overall, the results support alignment-based models of analogy and metaphor.

At the global level, the finding of metaphoric consistency effects across extended metaphor systems drives home the importance of systematicity in human analogical processing. These results are consistent with prior findings that people make analogical inferences based on higher-order connecting relations (Clement and Gentner 1991; Lassaline 1996; Markman 1997).

At present, only a handful of models of analogy and similarity incorporate the ability to use higher-order relations to constrain interpretation and inference (e.g., Falkenhainer, Forbus, and Gentner 1989; Holyoak and Thagard 1989; Keane and Brayshaw 1988). Of course, there are alternatives to higher-order relations. Relational connectivity can be modeled by defining groups of assertions that are processed together (Hummel and Holyoak 1997), and relational complexity can be modeled in terms of numbers of arguments to a relation, rather than the depth of the relational structure (Halford, Wilson, and Phillips 1998). However, our findings suggest that the phenomena of analogy and

metaphor—and even similarity—can best be captured by representing and using higher-order relations between relations.

Another link from metaphor to analogy is progressive abstraction. We have suggested that highlighting and storing the common schema is the chief mechanism by which novel metaphors become conventionalized. This process may also be central in learning ordinary categories. For example, the SEQL extension of SME, which simulates progressive abstraction of a common category (Skorstad, Gentner, and Medin 1988), has been used to simulate infants' regularity learning and adults' abstraction of story categories (Kuehne, Gentner, and Forbus 2000; Kuehne et al. 2000). We suggest that both metaphoric categories and literal categories (Kotovsky and Gentner 1996; Ross, Perkins, and Tenpenny 1990) can be derived by progressive abstraction across exemplars.

How Metaphors Are Different from Analogies

We have been stressing the commonalities between metaphor and analogy. There are also some differences. First, as we discussed at the outset, metaphors can be more structurally variable than analogies: They can be attribute matches, relation matches, or both; they can even violate structurally consistency (Gentner, Falkenhainer, and Skorstad 1988). As Fauconnier and Turner (1998) have noted, metaphors include complex blends that combine structure-mapping with metonymy and other processes. A second point is that the term *metaphor* is often used for novel and vivid nonliteral comparisons (Ortony 1975). For example, the subjects in Bowdle and Gentner's metaphor studies considered novel comparisons more metaphorical than conventional ones. However, as noted above, the term *metaphor* can also apply to systems of extended meanings that are so familiar as to be almost invisible (Lakoff and Johnson 1979; Nagy 1974).

Another dimension of difference is the pragmatic function of the figurative language. Gentner (1982) suggested that metaphors are typically used for expressive-affective purposes, and analogies for explanatory-predictive purposes. But we often speak of metaphors in science, so it might be more accurate to say that *analogy* is used in explanatory-predictive contexts, while *metaphor* can be used more broadly, in either explanatory-predictive or expressive-affective contexts.

From Comparison to Conceptual Systems: Metaphoric Systems in Reasoning

The process of conventionalization can result in stock metaphors, and finally in dead (or bleached, or frozen) local metaphorical senses. However, some metaphors take a different route, and end as conventionalized systems of reasoning. One case is which this occurs is in theory development in a domain. For example, in the history of theories of cognition, a common conceptualization of the mind is as a physical space (Roediger 1980). In this mapping, memories are objects in mental space, and recall involves spatial search. As cognitive theories evolved (and as technological advances created a greater set of potential bases), the set of metaphors enlarged. Gentner and Grudin (1985) surveyed the metaphors used in psychology over the ninety years since 1895 and traced a shift from general spatial metaphors to more complex systems metaphors—physics metaphors such *as associative force* and *goal gradient*, telephone switchboard metaphors, circuitry metaphors, and, eventually, computer metaphors. Boyd (1979) identified a number of terms derived from the "mind is a computer" metaphor, including "information processing," "encoding," "decoding," "indexing," "feedback," and "memory stores." (See also Fauconnier's discussion of the computer virus metaphor in chapter 7 of this volume.) Metaphors like this derive their force not from a local resemblance between physical objects and memory traces but rather from mapping the *system* of relationships in which these objects are embedded.

Some conventional analogical models in the physical domain also show systematic relational mappings despite conventionality. For example, Gentner and Gentner (1983) found -in the domain of electricity that people extended their analogical models of electricity—typically, either a *water flow* metaphor or a *moving crowd* metaphor—to reason about new circuits. Despite their conventionality, these folk analogical models retained some relational generativity (diSessa 1983).

Finally, space-time metaphors are a striking example of a conventional metaphor that retains a systematic generative structure (Bierwisch 1996; Boroditsky 2000; Clark 1973; Gentner forthcoming; Gentner and Imai 1992; Gentner, Imai, and Boroditsky in preparation; McGlone and Harding 1998; Traugott 1978). In many languages, including English,

there are two different metaphoric *space → time* systems: the *ego-moving* metaphor, wherein the observer's context progresses along the timeline toward the future, and the *time-moving* metaphor, wherein time is conceived of as a river or conveyor belt on which events are moving from the future to the past. To test whether these metaphors are psychologically processed as global systems, Gentner, Imai, and Boroditsky asked whether the metaphor-consistency effect found by Gentner and Boronat would occur here. Participants read statements about time, stated in terms of spatial metaphors—e.g., *Joe's birthday is approaching (Time-moving)* or *We are approaching the holidays* (*Ego-moving*)—andhad to indicate whether a specified event was past or future relative to the other. Consistent with Gentner and Boronat's studies, processing was slowed by a shift from one space-time metaphor to the other—evidence for their psychological coherence. Likewise, McGlone et al. found that participants interpreted an ambiguous statement—such as *Wednesday's meeting was moved one day forward*—as either Thursday or Tuesday, depending on whether the preceding metaphoric context was ego-moving or time-moving. As further evidence for the existence of space-time generative mappings, Boroditsky (2000) showed that hearing sentences about spatial relations primes analogous sentences in the time domain, but not the reverse.

Why do some metaphors, such as space-time metaphors, become conventionalized as systems, while others turn into local categories? One possibility is that metaphors evolve as into conventional systems when they continue to support new reasoning that requires concatenating relations. Thus one factor in the evolution of metaphorical systems might be the needs of the target domain: metaphors retain system properties if they are needed for close reasoning. A second factor is the nature of the base domain. The metaphors that evolve into conventional systems are often taken from space—which perhaps deserves "universal donor" status—or from other well understood, highly systematic base domains such as flowing water or device causality. These two factors are presumably connected: the intention to carry out sustained reasoning about the target favors the selection of a systematic base domain.

Summary

We have suggested that metaphor is like analogy—that the basic processes of analogy are at work in metaphor. Specifically, we suggest that structural alignment, inference projection, progressive abstraction, and re-representation are employed in the processing of metaphor. Viewing metaphor as analogy offers a unified account of many important phenomena and helps resolve some current conflicts.

We further propose that individual metaphors evolve over the course of their lives from alignment-based processing in the early stages to projection-based processing as they become conventionalized. Conventionalization often results in local metaphoric categories, but it can also take the form of large-scale conventional systems of metaphors. Finally, the ubiquitousness of metaphor demonstrates again our human capacity for seeing and using common relational patterns—in short, for analogical insight.

Acknowledgments

Please address correspondence and reprint requests to Dedre Gentner, Northwestern University, Department of Psychology, 2029 Sheridan Rd., Evanston, IL 60208-2710. This work was supported by ONR grant N00014-92-J-1098 and NSF grant SBR-9511757 to the first author.

This chapter was partially prepared while the first author was a Fellow at the Center for Advanced Study in the Behavioral Sciences. We are grateful for the financial support provided by the William T. Grant Foundation, award #95167795. We thank Ken Forbus, Arthur Markman, Ken Kurtz and Jeff Loewenstein for comments on this paper, and Kathleen Braun for help with the research and analyses.

Notes

1. Metaphors are often defined as nonliteral similarity comparisons. Metaphors are distinguished from similes by the inclusion of explicit comparison forms such as *like* in similes, but not metaphors. Thus examples 1, 4, and 5 are technically similes, and 2 and 3 are metaphors. However, the term *metaphor* is often used to encompass both forms. We will use the term *metaphor* in this broad sense and mark cases where it contrasts with *simile.*

2. To avoid confusion, we will use the term *metaphor* to refer to an individual figurative phrase, and *global metaphor* to refer to an extended mapping between two domains.

3. The metaphors used earlier were designed to be individually fairly novel (even though many of them belonged to known conceptual metaphors).

4. Local-to-global is not the same as bottom-up, a point that occasionally engenders confusion (e.g., Love, Rouder, and Wisniewski 1999). In SME, processing starts by identifying matching nodes at *any level* of the structure, from higher-order relations to concrete perceptual attributes. These local identities are then coalesced into global system-mappings (Falkenhainer, Forbus, and Gentner, 1989; Forbus, Gentner, and Law 1995).

5. We make the theoretical assumption that similarity of relational predicates can be expressed as partial identity. This avoids the circularity of defining similarity in terms of similarity. (If we define two things to be similar if their predicates are similar, this merely pushes the problem of defining similarity to the predicate level.) The partial-identity assumption is also psychologically advantageous in modeling the phenomenon of re-representation (e.g., Gentner et al. 1997; Gentner et al. 1995). Thus for two situations to be analogous means that they must have some set of identical relations.

6. The original SME exhaustively produced all possible interpretations, but this is psychologically implausible. Although the interpretations found by the greedy merge algorithm cannot be guaranteed to be maximal, the algorithm does very well. Forbus and Oblinger (1990) tested the greedy algorithm on a large set of analogies; on fifty-two out of fifty-six pairs, its top interpretation was identical to the best interpretation found in an exhaustive merge.

7. However, there was one exception: Gentner and Wolff (1997) found a base advantage for metaphors having highly conventional meanings and low (metaphorical) similarity (i.e., similarity of relations). No base advantage was found for novel metaphors, regardless of metaphorical similarity. This pattern led us to suggest a race between horizontal alignment (promoted by high similarity) and vertical matching with the base abstraction (promoted by high base conventionality).

8. The metaphor interference effects in these studies cannot be attributed to local lexical effects. Although there was no difference in processing time between forward and reversed metaphors, both required longer to reject than scrambled metaphors (that is, re-pairings of the terms from the metaphors).

9. In LISA, information is directionally projected from a driver to a recipient. This would seem to place it among the projection models. However, because LISA can shift between the two terms of an analogy as to which is driver and which is recipient, its processing predictions are not clear.

10. This requires modeling categorization as an alignment process between an abstract representation and a more concrete one. We think this may be a viable account, as discussed later.

11. This invites a further conjecture. If conventional metaphoric categories are processed as alignments, then what about standard categories? We speculate that categories in general are processed via structural alignment and mapping between abstract representations and concrete ones (see Ramscar and Pain 1996).

12. Interestingly, these directionality effects held only if the two passages were alignable. When the two items were not alignable (as independently rated), subjects showed no directional preference regardless of the relative coherence of the passages. This is consistent with our claim that analogical inference depends on alignment.

References

Allbritton, D. W, McKoon, G., and Gerrig, R. (1995). Metaphor-based schemas and text representations: Making connections through conceptual metaphors. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 21:1–4.

Bierwisch, M. (1996). How much space gets into language? In P. Bloom, M. A. Peterson, L. Nadel, and M. F. Garrett, Eds., *Language and space*, pp. 31-76. Cambridge, MA: MIT Press.

Black, M. (1962). Metaphor. In M. Black, Ed., *Models and metaphors*, pp. 25–47. Ithaca, NY: Cornell University Press.

Blank, G. D. (1988). Metaphors in the lexicon. *Metaphor and Symbolic Activity* 3:21-26.

Blasko, D. G., and Connine, C. M. (1993). Effects of familiarity and aptness on metaphor processing. *Journal of Experimental Psychology: Learning, Memory and Cognition* 12:205-308.

Boroditsky, L. (2000). Metaphoric structuring: Understanding time through spatial metaphors. *Cognition* 75(1): 1-27.

Boronat, C. B. (1990). *Effects of base shift and frequency in extended metaphor processing.* M.A. thesis, University of Illinois at Urbana-Champaign.

Boronat, C., and Gentner, D. (in preparation). Novel metaphors are processed as generative domain-mappings.

Bowdle, B., and Gentner, D. (1995). The career of metaphor. Paper presented at the meeting of the Psychonomics Society, Los Angeles, CA.

Bowdle, B., and Gentner, D. (1997). Informativity and asymmetry in comparisons. *Cognitive Psychology* 34(3):244–286.

Bowdle, B., and Gentner, D. (1999). Metaphor comprehension: From comparison to categorization. *Proceedings of the twenty-first annual conference of the Cognitive Science Society.* Hillsdale, NJ: Erlbaum.

Bowdle, B., and Gentner, D. (in preparation). *The career of metaphor.*

Boyd, R. (1979). Metaphor and theory change: What is "metaphor" a metaphor for? In A. Ortony, Ed., *Metaphor and thought*, pp. 356–408. Cambridge: Cambridge University Press.

Burstein, M. H. (1983). Concept formation by incremental analogical reasoning and debugging. *Proceedings of the International Machine Learning Workshop,* 19-25.

Burstein, M. H. (1986). Concept formation by incremental analogical reasoning and debugging. In R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, Eds., *Machine learning: An artificial intelligence approach,* vol. 2, pp. 351-369. Los Altos, CA: Kaufmann.

Cacciari, C., and Tabossi, P. (1988). The comprehension of idioms. *Journal of Memory and Language* 27:668-683.

Camac, M. K., and Glucksberg, S. (1984). Metaphors do not use associations between concepts, they are used to create them. *Journal of Psycholinguistic Research* 13:443–455.

Clark, H. H. (1973). Space, time, semantics, and the child. In T. E. Moore, Ed., *Cognitive development and the acquisition of language,* pp. 27-63. New York: Academic Press.

Clark, H. H., and Lucy, P. (1975). Understanding what is meant from what is said: A study in conversationally conveyed requests. *Journal of Verbal Learning and Verbal Behavior* 14:56-72.

Clement, C. A., and Gentner, D. (1991). Systematicity as a selection constraint in analogical mapping. *Cognitive Science* 15:89-132.

diSessa, A. A. (1983). Phenomenology and the evolution of intuition. In D. Gentner and A. L. Stevens, Eds., *Mental models,* pp. 15-33. Hillsdale, NJ: Erlbaum.

Falkenhainer, B., Forbus, K. D., and Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence* 41:1-63.

Fauconnier, G. (1990). Domains and connections. *Cognitive Linguistics* 1-1:151-174.

Fauconnier, G., and Turner, M. (1998). Conceptual integration networks. *Cognitive Science* 22(2):133–187.

Forbus, K. D., Ferguson, R. W., and Gentner, D. (1994). Incremental structure-mapping. *Proceedings of the sixteenth annual conference of the Cognitive Science Society,* pp. 313-318. Mahwah, NJ: Erlbaum.

Forbus, K. D., and Gentner, D. (1989). Structural evaluation of analogies: What counts? *Proceedings of the eleventh annual conference of the Cognitive Science Society,* pp. 341-348. Mahwah, NJ: Erlbaum.

Forbus, K. D., Gentner, D., and Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science* 19(2):141-205.

Forbus, K. D., and Oblinger, D. (1990). Making SME greedy and pragmatic. *Proceedings of the twelfth annual conference of the Cognitive Science Society,* pp. 61-68. Mahwah, NJ: Erlbaum.

Gentner, D. (1975). Evidence for the psychological reality of semantic components: The verbs of possession. In D. A. Norman, D. E. Rumelhart, and the LNR Research Group, Eds., *Explorations in cognition, pp.* 211-246. San Francisco: W.H. Freeman.

Gentner, D. (1982). Are scientific analogies metaphors? In D. Miall, Ed., *Metaphor: Problems and perspectives*, pp. 106-132. Brighton: Harvester.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science* 7:155-170.

Gentner, D. (1988). Metaphor as structure mapping: The relational shift. *Child Development* 59:47-59.

Gentner, D. (1992). Metaphor as mapping. Paper presented at the meeting of the Cognitive Science Society, Chicago.

Gentner, D. (forthcoming). Spatial metaphors in temporal reasoning. In M. Gattis, Ed., *Spatial schemas in abstract thought.* Cambridge, MA: MIT Press.

Gentner, D., and Boronat, C. B. (1992). Metaphor as mapping. Paper presented at the Workshop on Metaphor, Tel Aviv.

Gentner, D., and Boronat, C. (in preparation). Metaphors are (sometimes) processed as generative domain-mappings.

Gentner, D., and Bowdle, B. (1994). The coherence imbalance hypothesis: A functional approach to asymmetry in comparison. *Proceedings of the sixteenth annual meeting of the Cognitive Science Society,* 351-356. Hillsdale, NJ: Erlbaum.

Gentner, D., Brem, S., Ferguson, R. W., Markman, A. B., Levidow, B. B., Wolff, P., and Forbus, K. D. (1997). Analogical reasoning and conceptual change: A case study of Johannes Kepler. *The Journal of the Learning Sciences* 6(1):3–40.

Gentner, D., and Clement, C. A. (1988). Evidence for relational selectivity in the interpretation of analogy and metaphor. In G. H. Bower, Ed., *The psychology of learning and motivation,* pp. 307-358. New York: Academic.

Gentner, D., Falkenhainer, B., and Skorstad, J. (1988). Viewing metaphor as analogy. In D. H. Helman, Ed., *Analogical reasoning: Perspectives of artificial intelligence, cognitive science, and philosophy,* pp. 171-177. Dordrecht: Kluwer.

Gentner, D., and France, I. M. (1988). The verb mutability effect: Studies of the combinatorial semantics of nouns and verbs. In S. L. Small, G. W. Cottrell, and M. K. Tanenhaus, Eds., *Lexical ambiguity resolution: Perspectives from psycholinguistics, neuropsychology, and artificial intelligence,* pp. 343-382. San Mateo, CA: Kaufmann.

Gentner, D., and Gentner, D. R. (1983). Flowing waters or teeming crowds: Mental models of electricity. In D. Gentner and A. L. Stevens, Eds., *Mental models,* pp. 99-129. Hillsdale, NJ: Erlbaum.

Gentner, D., and Grudin, J. (1985). The evolution of mental metaphors in psychology: A 90-year retrospective. *American Psychologist* 40:181-192.

Gentner, D., and Imai, M. (1992). Is the future always ahead? Evidence for system-mappings in understanding space-time metaphors. *Proceedings of the fourteenth annual meeting of the Cognitive Science Society,* pp. 510-515. Mahwah, NJ: Erlbaum.

Gentner, D., Imai, M., and Boroditsky, L. (in preparation). As time goes by: Evidence for two systems in processing space-time metaphors.

Gentner, D., and Jeziorski, M. (1993). The shift from metaphor to analogy in western science. In A. Ortony, Ed., *Metaphor and thought* (2d ed.), pp. 447–480. Cambridge: Cambridge University Press.

Gentner, D., and Markman, A. B. (1997). Structure mapping in analogy and similarity. *American Psychologist* 52:45-56.

Gentner, D., Rattermann, M. J., and Forbus, K. D. (1993). The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology* 25:524-575.

Gentner, D., Rattermann, M. J., Markman, A. B., and Kotovsky, L. (1995). Two forces in the development of relational similarity. In T. J. Simon and G. S. Halford, Eds., *Developing cognitive competence: New approaches to process modeling*, pp. 263-313. Hillsdale, NJ: Erlbaum.

Gentner, D., and Wolff, P. (2000). Metaphor and knowledge change. In E. Dietrich and A. B. Markman, Eds., *Cognitive dynamics: Conceptual and representational change in humans and machines,* pp. 295-342. Mahwah, NJ: Erlbaum.

Gentner, D., and Wolff, P. (1997). Alignment in the processing of metaphor. *Journal of Memory and Language* 37:331-355.

Gerrig, R. J., and Gibbs, R. W. (1988). Beyond the lexicon: Creativity in language production. *Metaphor and Symbolic Activity* 3:1–19.

Gibbs, R. (1985). On the process of understanding idioms. *Journal of Psycholinguistic Research* 14:465–472.

Gibbs, R. W, Jr. (1979). Contextual effects in understanding indirect requests. *Discourse Processes* 2:149-156.

Gibbs, R. W, Jr. (1980). Spilling the beans on understanding and memory for idioms in conversations. *Memory and Cognition* 8:449–456.

Gibbs, R. W, Jr. (1990). Comprehending figurative referential descriptions. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 16:56-66.

Gibbs, R. W, Jr. (1994). *The poetics of mind: Figurative thought, language, and understanding.* New York: Cambridge University Press.

Gibbs, R. W, Jr. (1996). Why many concepts are metaphorical. *Cognition* 61:309-319.

Gibbs, R., Nayak, N. P., and Cutting, C. (1989). How to kick the bucket and not decompose: Analyzability and idiom processing. *Journal of Memory and Language* 28:576-593.

Gibbs, R., and O'Brien, J. E. (1990). Idioms and mental imagery: The metaphorical motivation for idiomatic meaning. *Cognition* 36(1):35-68.

Gick, M. L., and Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology* 15:1–38.

Gildea, P., and Glucksberg, S. (1983). On understanding metaphors: The role of context. *Journal of Verbal Learning and Verbal Behavior* 22:577-590.

Giora, R. (1997). Understanding figurative and literal language: The graded salience hypothesis. *Cognitive Linguistics* 8(3):183–206.

Glucksberg, S., Brown, M., and McGlone, M. S. (1993). Conceptual analogies are not automatically accessed during idiom comprehension. *Memory and Cognition* 21:711–719.

Glucksberg, S., Gildea, P., and Bookin, H. B. (1982). On understanding nonliteral speech: Can people ignore metaphors? *Journal of Verbal Learning and Verbal Behavior* 21:85-98.

Glucksberg, S., and Keysar, B. (1990). Understanding metaphorical comparisons: Beyond similarity. *Psychological Review* 97:3-18.

Glucksberg, S., McGlone, M. S., and Manfredi, D. (1997). Property attribution in metaphor comprehension. *Journal of Memory and Language* 36:50-67.

Goldstone, R. L. (1994). Similarity, interactive activation, and mapping. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 20(1):3–28.

Goldstone, R. L., and Medin, D. L. (1994). Similarity, interactive-activation and mapping: An overview. In K. J. Holyoak and J. A. Barnden, Eds., *Advances in connectionist and neural computation theory*, vol. 2, *Analogical connections*, pp. 321-362. Norwood, NJ: Ablex.

Greiner, R. (1988). Learning by understanding analogies. *Artificial Intelligence* 35:81-125.

Gruber, H. E. (1995). Insight and affect in the history of science. In R. J. Sternberg and J. E. Davidson, Eds., *The nature of insight*, pp. 397–432. Cambridge, MA: MIT Press.

Halford, G. S., Wilson, W. H., and Phillips, S. (1998). Processing capacity defined by relational complexity: Implications for comparative, developmental, and cognitive psychology. *Behavioral and Brain Sciences* 21:803-864.

Hobbs, J. R. (1979). Metaphor, metaphor schemata, and selective inferencing. Technical Note 204, SRI Projects 7910 and 7500. Menlo Park, CA: SRI International.

Hoffman, R. R., and Kemper, S. (1987). What could reaction-time studies be telling us about metaphor comprehension? *Metaphor and Symbolic Activity* 2:149-186.

Holyoak, K. J., and Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science* 13(3):295-355.

Holyoak, K. J., and Thagard, P. R. (1995). *Mental leaps: Analogy in creative thought.* Cambridge, MA: MIT Press.

Honeck, R. P., Kibler, C. T., and Firment, M. J. (1987). Figurative language and psychological views of categorization: Two ships in the night? In R. E. Haskell, Ed., *Cognition and symbolic structures: The psychology of metaphoric transformation*, pp. 103-120. Norwood, NJ: Ablex.

Hummel, J. E., and Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review* 104(3):427–466.

Keane, M. T., and Brayshaw, M. (1988). The incremental analogical machine: A computational model of analogy. In D. Sleeman, Ed., *Third European working session on machine learning*, pp. 53-62. San Mateo, CA: Kaufmann.

Keane, M. T, Ledgeway, T, and Duff, S. (1994). Constraints on analogical mapping: A comparison of three models. *Cognitive Science* 18(3):387–438.

Kedar-Cabelli, S. (1988). Toward a computational model of purpose-directed analogy. In A. Prieditis, Ed., *Analogica*, pp. 89-107. Los Altos, CA: Kaufmann.

Kelly, M. H., and Keil, F. C. (1987). Metaphor comprehension and knowledge of semantic domains. *Metaphor and Symbolic Activity* 2:33-51.

Kennedy, J. M. (1990). Metaphor—Its intellectual basis. *Metaphor and Symbolic Activity* 5:115-123.

Keysar, B. (1989). On the functional equivalence of literal and metaphorical interpretations in discourse. *Journal of Memory and Language* 28:375–385.

Keysar, B., and Bly, B. (1995). Intuitions of the transparency of idioms: Can one keep a secret by spilling the beans? *Journal of Memory and Language* 34(1):89-109.

Kintsch, W. (1974). *The representation of meaning in memory.* Hillsdale, NJ: Erlbaum.

Kittay, E. F. (1987). *Metaphor: Its cognitive force and linguistic structure.* Oxford: Clarendon.

Kittay, E. E, and Lehrer, A. (1981). Semantic fields and the structure of metaphor. *Studies in Language* 5(l):31-63.

Kotovsky, L., and Gentner, D. (1996). Comparison and categorization in the development of relational similarity. *Child Development* 67:2797-2822.

Kuehne, S. E., Forbus, K. D., Gentner, D., and Quinn, B. (2000). SEQL: Category learning as incremental abstraction using structure mapping. In *Proceedings of the twenty-second annual conference of the Cognitive Science Society.* Hillsdale, NJ: Erlbaum.

Kuehne, S. E., Gentner, D., and Forbus, K. D. (2000). Modeling infant learning via symbolic structural alignment. *Proceedings of the twenty-second annual Conference of the cognitive science society.* Hillsdale, NJ: Erlbaum Associates.

Lakoff, G. (1990). The invariance hypothesis: Is abstract reason based on image-schemas? *Cognitive Linguistics* l(l):39-74.

Lakoff, G. (1993). The contemporary theory of metaphor. In A. Ortony, Ed., *Metaphor and thought*, 2nd ed., pp. 202-251. New York: Cambridge University Press.

Lakoff, G., and Johnson, M. (1980). *Metaphors we live by*. Chicago: University of Chicago Press.

Lakoff, G., and Turner, M. (1989). *More than cool reason: A field guide to poetic metaphor.* Chicago: University of Chicago Press.

Lassaline, M. E. (1996). Structural alignment in induction and similarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22(3):754-770.

Lehrer, A. (1990). Polysemy, conventionality, and the structure of the lexicon. *Cognitive Linguistics* l(2):207-246.

Loewenstein, J., Thompson, L., and Gentner, D. (1999). Analogical encoding facilitates knowledge transfer in negotiation. *Psychonomic Bulletin and Review* 6(4):586-597.

Love, B. C., Rouder, J. N., and Wisniewski, E. J. (1999). A structural account of global and local processing. *Cognitive Psychology* 38:291-316.

Malgady, R., and Johnson, M. (1980). Measurement of figurative language: Semantic feature models of comprehension and appreciation. In R. Honeck and R. Hoffman, Eds., *Cognition and figurative language*, pp. 239-258. Hillsdale, NJ: Erlbaum.

Markman, A. B. (1997). Constraints on analogical inference. *Cognitive Science* 21(4):373–418.

Markman, A. B. (1999). *Knowledge representation.* Mahwah, NJ: Erlbaum.

Markman, A. B., and Gentner, D. (1993). Structural alignment during similarity comparisons. *Cognitive Psychology* 25:431-467.

Marschack, M., Katz, A., and Paivio, A. (1983). Dimensions of metaphor. *Journal of Psycholinguistic Research* 12:17-40.

McGlone, M. S. (forthcoming). Concepts as metaphors. In S. Glucksburg, Ed., *Metaphor and allusion: Studies of figurative language comprehension.* Oxford: Oxford University Press.

McGlone, M. S., and Harding, J. (1998). Back (or forward?) to the future: The role of perspective in temporal language comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 24(5):1211–1223.

Medin, D. L., Goldstone, R. L., and Gentner, D. (1993). Respects for similarity. *Psychological Review* 100(2):254-278.

Miller, G. A. (1979). Images and models, similes and metaphors. In A. Ortony, Ed., *Metaphor and thought*, pp. 202-250. Cambridge: Cambridge University Press.

Miller, G. A. (1993). Images and models, similes and metaphors. In A. Ortony, Ed., *Metaphor and thought* (2d ed.), pp. 357–400. Cambridge: Cambridge University Press.

Murphy, G. L. (1996). On metaphoric representation. *Cognition* 60(2):173-204.

Murphy, G. L. (1997). Reasons to doubt the present evidence for metaphoric representation. *Cognition* 62:99-108.

Nagy, W. (1974). *Figurative patterns and redundancy in the lexicon.* Ph.D. dissertation, University of California at San Diego.

Nayak, N. P., and Gibbs, R. W. (1990). Conceptual knowledge in the interpretation of idioms. *Journal of Experimental Psychology* 119(3):315-330.

Nersessian, N. J. (1992). How do scientists think? Capturing the dynamics of conceptual change in science. In R. N. Giere and H. Feigl, Eds., *Cognitive models of science: Minnesota studies in the philosophy of science*, pp. 3–44. Minneapolis: University of Minnesota Press.

Novick, L. R., and Holyoak, K. J. (1991). Mathematical problem solving by analogy. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 17(3):398–415.

Ortony, A. (1975). Why metaphors are necessary and not just nice. *Educational Theory* 25:45-53.

Ortony, A. (1979). Beyond literal similarity. *Psychological Review* 86:161-180.

Ortony, A., Schallert, D. L., Reynolds, R. E., and Antos, S. J. (1978). Interpreting metaphors and idioms: Some effects of context on comprehension. *Journal of Verbal Learning and Verbal Behavior* 17:465–477.

Ortony, A., Vondruska, R. J., Foss, M. A., and Jones, L. E. (1985). Salience, similies, and the asymmetry of similarity. *Journal of Memory and Language* 24:569-594.

Quinn, N. (1987). Convergent evidence for a cultural model of American marriage. In D. Holland and N. Quinn, Eds., *Cultural models in language and thought*, pp. 173-192. New York: Cambridge University Press.

Ramscar, M., and Pain, H. (1996). Can a real distinction be made between cognitive theories of analogy and categorisation? *Proceedings of the eighteenth annual conference of the Cognitive Science Society*, pp. 346-351. Mahwah, NJ: Erlbaum.

Roediger, H. (1980). Memory metaphors in cognitive psychology. *Memory and Cognition* 8:231–246.

Ross, B. H., and Kennedy, P. T. (1990). Generalizing from the use of earlier examples in problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 16:42-55.

Ross, B. H., Perkins, S. J., and Tenpenny, P. L. (1990). Reminding-based category learning. *Cognitive Psychology* 22:460–492.

Rumelhart, D. E., and Abrahamson, A. A. (1973). A model for analogical reasoning. *Cognitive Psychology* 5:1-28.

Searle, J. R. (1976). A classification of illocutionary acts. *Language in Society* 5(1):1–23.

Shen, Y. (1992). Metaphors and categories. *Poetics Today* 13:771-794.

Skorstad, J., Gentner, D., and Medin, D. (1988). Abstraction processes during concept learning: A structural view. *Proceedings of the tenth annual conference of the Cognitive Science Society*, pp. 419–425. Mahwah NJ: Erlbaum.

Steen, G. J. (1989). Metaphor and literary comprehension: Towards a discourse theory of metaphor in literature. *Poetics*, 18:113-141.

Sweetser, E. (1990). *From etymology to pragmatics.* Cambridge: Cambridge University Press.

Swinney, D., and Cutler, A. (1979). The access and processing of idiomatic expressions. *Journal of Verbal Learning and Verbal Behavior* 18:523-534.

Tourangeau, R., and Rips, L. (1991). Interpreting and evaluating metaphors. *Journal of Memory and Language* 30:452–472.

Tourangeau, R., and Sternberg, R. (1981). Aptness in metaphor. *Cognitive Psychology* 13:27-55.

Tourangeau, R., and Sternberg, R. (1982). Understanding and appreciating metaphors. *Cognition* 11:203-244.

Traugott, E. C. (1978). On the expression of spatio-temporal relations in language. In J. H. Greenberg, Ed., *Universals of human language*, vol. 3, *Word structure*, pp. 369–400. Stanford: Stanford University Press.

Turner, M. (1987). *Death is the mother of beauty: Mind, metaphor, and criticism.* Chicago: University of Chicago Press.

Turner, M. (1991). *Reading minds: The study of English in the age of cognitive science.* Princeton, NJ: Princeton University Press.

Turner, N. E., and Katz, A. N. (1997). The availability of conventional and of literal meaning during the comprehension of proverbs. *Pragmatics and Cognition* 5:199-233.

Williams, J. (1992). Processing polysemous words in context: Evidence for interrelated meanings. *Journal of Psycholinguistic Research* 21:193-218.

Wolff, P., and Gentner, D. (1992). The time course of metaphor comprehension. *Proceedings of the fourteenth annual conference of the Cognitive Science Society,* pp. 504-509. Mahwah, NJ: Erlbaum.

Wolff, P., and Gentner, D. (2000). Evidence for role-neutral initial processing of metaphors. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 26(2):529-541.

Wolff, P., and Gentner, D. (in preparation). From symmetric to asymmetric processing: Two stages in the comprehension of metaphors.