How one's hook is baited matters for catching an analogy

Jeffrey Loewenstein

The University of Texas at Austin

1 University Station B6300

Austin, TX 78712

512-475-8020

512-471-3937 (fax)

jeffrey.loewenstein@mccombs.utexas.edu

Contents

Abstract

Memory provides an ocean of possibilities, so it is necessary to find good ways to bait one's hook to ensure catching something worthwhile. Intelligent action requires linking useful, previously learned examples with current problems, which very often means intelligent action requires retrieving analogies. We already know that effectively encoding the underlying structure in examples during initial learning facilitates later retrieving them as analogies. My colleagues and I have recently found that it is also possible to facilitate retrieving analogies by effectively encoding the example serving as a probe to memory, without relying on any special encoding of the stored examples. Thus people do not have to learn examples well initially to still make good use of them, a finding with useful implications for problem solving, creativity, and acquiring domain knowledge.

1. Introduction

Acting intelligently requires information, leading to the fundamental problem of having the right information at the right time. Perhaps the key challenge making this problem difficult is readily retrieving a small bit of relevant information from a vast amount of available information. As an example of the scale and importance of the retrieval challenge, in a mundane search of the web, people usually look at no more than a handful of the search results retrieved; those results are drawn from a pool of over a trillion pages on the web (as of July 2008); and the market value of the web search industry is on the scale of hundreds of billions of dollars. Compared to retrieving web pages, retrieving information from one's own memory is a more challenging and more critical problem. So, it is important to study the challenges people have in retrieving information from memory (Gentner, Rattermann & Forbus, 1993; Ross, 1984). It is also why my colleagues and I (Gentner, Loewenstein, Thompson & Forbus, 2009; Kurtz & Loewenstein, 2007) have been examining whether there is a new avenue for retrieving information effectively.

Retrieving the right information at the right time is only sometimes challenging for people. Often, people do so effortlessly, so we consider the retrievals mundane—remembering what a dog is or how to get dressed are not typically recognized as achievements. More generally, the world is often kind (Gentner, 1989; Gentner & Medina, 1998; Medin & Ross, 1989), in the sense that surface

appearances are often correlated with the underlying structures, principles or essences required to act intelligently. This allows people a short cut. They can encode surface properties and retrieve information from memory based on surface properties. Of course, one cannot always judge a book by its cover. It is in these sorts of situations—when information that is relevant comes from a source without surface properties in common—that people find it challenging to retrieve information and as a result fail to act intelligently. These are situations in which people need to retrieve examples that have different surface properties but the same underlying structures. People need to retrieve *analogies* (Gentner, 1983). Thus, the biggest challenge of people retrieving the right information at the right time from their own memories is the challenge of how to retrieve analogies.

The dominant trend in research on retrieving analogies has been to emphasize the importance of how information is initially indexed or encoded (e.g., Schank, 1982; Gick & Holyoak, 1983). It is reasonable that for information to be later retrieved at the right times, people need to encode it well from the start. It also makes sense in an educational context to consider how people can learn information well initially such that they are likely to be able to use it later when it is relevant. A less obvious reason to focus on initial encoding is that in addition to learning new examples, initial encoding is the point at which people can also be learning new indexing terms (or more broadly, learning new encoding vocabularies; Forbus, Gentner & Law, 1995; Kolodner, 1993). If people learn new indexing terms and then use those terms to encode future examples, this should facilitate retrieving the initial examples encoded with the same indexing terms. This "learning to encode"

account (Gentner et al., 2009), together with the educational focus and the interest in people encoding information well from the start, provide reasons for focusing on how information is initially indexed when considering the retrieval of analogies.

The largely unasked question has been whether improving people's initial encoding of examples is the only influence on retrieving analogies. This is important, because if changing initial encoding is the only means for influencing the retrieval of analogies, then we have little to say to someone currently confronted by a challenge in retrieving an analogy. Their prior knowledge is either already well encoded (and hence retrieving analogies should not be so challenging) or it is effectively useless because it cannot be retrieved. But it appears that there is another means for facilitating the retrieval of analogies. My colleagues and I (Gentner et al., 2009; Kurtz & Loewenstein, 2007) have recently documented that people can change their ability to retrieve analogies at the time of retrieval, without changing initial learning.[1]

2. Key roles for retrieving analogies

The difficulties posed by retrieving analogies appear as specific challenges in many areas, consistent with effective information retrieval being a fundamental problem. Retrieving analogies is the central challenge behind effective knowledge transfer during problem solving (Gick & Holyoak, 1980; Ross, 1989). Retrieving

---

[1] As will be apparent, this stream of research and my thinking on it has benefitted greatly from the influence of my mentor and collaborator, Dedre Gentner; the errors are mine alone.

analogies is a central challenge for creativity, as people need to connect new kinds of

relevant information to current situations to produce novel and useful outcomes

(Markman, Wood, Linsey, Murphy & Laux, 2009). Retrieving analogies is also a key

challenge in the process of acquiring domain knowledge, both in the course of child

development (Brown, 1989, 1990) and in the course of developing expertise (Chi &

Ohlsson, 2005). This is because people need to learn, apply and organize their

knowledge using the underlying principles in the domain rather than surface

properties. Consequently, before considering a new way in which people can

retrieve analogies, I examine prior research on retrieving analogies in problem

solving, creativity, and the acquisition of domain knowledge.

2.1   Problem solving and retrieving analogies

Research on problem solving has long been concerned with people's ability

(or lack of ability) to transfer knowledge from one problem or situation to help

solve further problems (Reeves & Weisberg, 1994; Ross, 1984; Whitehead, 1929).

For example, people negotiating the syndication rights to a television show might

recall an earlier negotiation in which the cost of shipping was contingent on the

timing of the goods arrival, and realize that they could make the payments for the

television show contingent on the show's popularity (Thompson, Gentner &

Loewenstein, 2000). Spontaneously retrieving a prior example because it has

matching underlying structure (e.g., contingent payments) rather than matching

surface features (e.g., negotiating over a television show) seems to be the main

sticking point. This is because people working on a problem who fail to retrieve a

prior analogous example spontaneously can very often succeed at making use of the

earlier example if given a hint to use it (Gick & Holyoak, 1980). Even without the

potential distraction of working on a problem, people who read an example and are

asked to recall an earlier example typically fail to retrieve an analogy, but can select

the appropriate analogical match if given a forced choice (Gentner, Rattermann &

Forbus, 1993). Accordingly, how to apply a prior example to help solve a current

problem is usually less of a challenge (but still a concern, e.g., Bassok & Holyoak,

1989; Blessing & Ross, 1996) than the challenge of retrieving the prior example.

Research provides varying degrees of support for factors that enable people

to retrieve analogies during problem solving. There is widespread agreement that

having drawn a comparison between initial examples facilitates retrieving them

later when confronted by an analogous problem or situation. This effect holds when

people compare initial examples as part of a learning phase (e.g., Catrambone &

Holyoak, 1989; Gentner, Loewenstein & Thompson, 2003; Gick & Holyoak, 1983)

and when people compare initial examples in the course of problem solving (e.g.,

Dixon & Bangert, 2004; Ross & Kennedy, 1990; see also Gentner, Loewenstein &

Hung, 2007; Loewenstein & Gentner, 2001). There is also evidence from a variety of

sources that linking an example not to a second example but to an abstract principle

(a description of the underlying structure in an example) can promote later

analogical transfer (Goldstone & Wilensky, 2008). This effect has been found, for

example, for people who read a principle embedded within an example (Ross &

Kilbane, 1997), people who have to select which of a small number of principles

applies to an example (Seifert, McKoon, Abelson & Ratcliff, 1986), and people who

compare a principle and an example (Gentner, Loewenstein & Thompson, 2004).

Related work, although not showing effects on problem solving, does show that people with sufficient relevant background knowledge and encouragement to study single examples can derive the underlying structure in those examples through self-explanation (Ahn, Brewer & Mooney, 1992) or discussion (Schwartz, 1995). It is plausible that self-explanation or joint discussion could therefore facilitate later analogical retrieval of the studied example and promote problem solving. There is a small amount of evidence that even without an explanation, just using a few words to emphasize the underlying structure in the initial example facilitates using that example later: Loewenstein and Gentner (2005, Experiment 5) found transfer effects between mapping tasks, and Clement, Mawby, and Giles (1994) found effects on retrieving analogies. An alternative approach to emphasizing underlying structure is to remove potential distractions from examples. There is a body of literature supporting the idea that both using sparse examples with few surface features (DeLoache, 1995; Gentner & Rattermann, 1991) and progressively removing concrete details to make an example more schematic (Goldstone & Son, 2005) promote later transfer to solving an analogous problem. Finally, a new finding is that presenting written examples, as is commonly done in studies of analogical retrieval and problem solving, engenders lower rates of analogical retrieval than presenting those examples in spoken form (Markman, Taylor & Gentner, 2007). It is possible that the advantage of hearing rather than reading examples is due to listening encouraging people to emphasize the gist of the examples.

A potential generalization across all these findings is that people need encouragement to generate encodings of examples that (1) articulate the underlying

structure in examples, (2) articulate that underlying structure in a generic fashion, apart from contextual details and (3) emphasize that underlying structure. It cannot simply be that people fail to encode underlying structure and these interventions help them to do so—the primary stumbling block is retrieving the analogy, not using the analogy once it is retrieved. If people did not encode the information, it would not be available for application. Thus the effects of comparing examples, linking examples to abstract principles, self-explanations of the underlying structure in examples, language for the underlying structure in examples and fading away details must be more about how and how much of the underlying structure is encoded.

The various interventions seem to encourage people to encode the complete underlying structure in examples and to do so in a generic way. A generalization or schema resulting from comparison should capture all the commonalities between examples, and because it has to apply to both examples it must be at least somewhat generic—how generic will depend on how different the examples are. Abstract principles are necessarily generic. A component of self-explanations is developing and improving one's understandings of the underlying structure in examples. Language for the underlying structure in examples can provide generic categories for characterizing the key underlying structure in examples and can point out critical aspects of that structure thereby minimizing the likelihood of partial understandings. Fading away details also makes examples less specific and more generic, as well as enhances the ease of perceiving the full underlying structure. Finally, gist encodings of examples should also tend towards the generic. To be

clear, all of these routes to fostering complete and generic encodings of examples are efforts at removing details and preserving the underlying structure in the examples. They are not fostering the derivation of vapid abstractions ("something happened") that gloss over the coherent systems of roles and relations in the examples.

The final commonality to the interventions is that they foster emphasizing the underlying structure in the examples. This is not to say that people fail to encode other aspects of examples. For example, we know that people continue to use surface features in the retrieval and application of prior examples even after they retrieve analogies (Bassok & Holyoak, 1989; Blessing & Ross, 1996; Novick, 1988). It is a matter of pushing that surface property information to the background and bringing underlying structure to the fore.

One way to think about the interventions is that by fostering an emphasis on complete, generic encodings of underlying structure in examples, they are encouraging people to encode examples like domain experts. Presumably experts focus their encodings on the underlying structure in examples, consistent with research that experts organize problems based on their underlying principles rather than their surface properties (e.g., Chi, Feltovich & Glaser, 1981). If so, novices should have the most difficulty retrieving analogies, whereas experts should have less difficulty retrieving analogies. In support of this claim, Novick (1988) found that students with higher mathematics test scores were better able than those with lower test scores to retrieve analogies to solve problems. Also in support of this account, Dunbar and Blanchette (2001) report that microbiology researchers (who

11

were presumably experts in the domain of microbiology) generated more spontaneous analogies between areas of microbiology than they did to examples outside the domain.

Also consistent with this account, my colleagues and I (Gentner et al, 2009) have found that higher base rates for retrieving analogies from one's own memory were associated with more years of domain experience. I will provide specifics on our data, and supplement the published data with unpublished data, because our work extends prior research on the relationship between experience and analogical retrieval. Collapsing across several studies that used the same methodology, we have found that the quality (i.e., degree of structural match, measured on a 0 to 2 scale) of people's retrievals based on a negotiation example was associated with the extent of their business experience. We found low quality retrievals from those with no business experience (M = 0.24, n = 17), better quality retrievals from people with one to seven years of business experience (M = 0.52, n = 81), and still higher quality retrievals from people with 15 to 40 years of business experience (M = 0.84, n = 80), χ^2 (2, N=178) = 10.81, p < 0.05. This is weak evidence given that it consists of cross-study comparisons, but it is suggestive of the general pattern claimed in the literature and covers a more substantial range of experience than previous research.

Taken together, the research on retrieving analogies in problem solving can largely be summarized by claiming that better encodings of examples today lead to better retrievals of those examples tomorrow. If people are encouraged to encode examples in more sophisticated ways than they apparently normally do—such as

because drew analogical comparisons—then they are more likely to be able to later retrieve the example when they encounter an analogous problem.

2.2   Creativity and retrieving analogies

Research on creativity has long emphasized both the value of analogies (Finke, Ward & Smith, 1992; Hesse, 1966; Hofstadter & FARG, 1995; Koestler 1969; Weisberg, 1993) and the value of connecting ideas from different domains (Mednick, 1962) as a way to generate novel and useful outcomes. However, creativity research has not always linked these two discussions, perhaps because creativity research is interested in all useful associations, and even random associations (Campbell, 1960). Still, analogies provide an important avenue for the generation of creative outcomes (Holyoak & Thagard, 1995), and accordingly analogical retrieval is important for creativity and innovation (Markman, Wood, Linsey, Murphy & Laux, 2009). For example, a topic of strong current interest within the design and innovation community is biomimicry (e.g., biomimicry.net; asknature.org), or the generation of designs based on analogies to nature. Benyus (1997) describes the modification of a train's design so that it entered tunnels more efficiently based on an analogy to how kingfisher birds enter water. As this example indicates, retrieving an appropriate analogy to support creativity can be daunting, as it was within the domain of problem solving. Perhaps this is unsurprising, given that some researchers treat creativity as a kind of problem solving (e.g., Amabile, 1996). Still, creativity research has developed along somewhat different lines than problem solving research, and provides some distinct insights into the retrieval of analogies.

One area of creativity research has examined providing various sorts of memory cues and their influence on analogical retrieval. This is related to the research discussed earlier on the world being kind, which suggested that surface properties can serve as the primary basis for retrievals that can subsequently be used as analogies (Blessing & Ross, 1996; Ross, 1987). Research on brainstorming (Dugosh et al, 2000) and on design tasks (Christensen & Schunn, 2009) suggests that providing multiple retrieval cues as people work on creative tasks facilitates their performance, and at least some of these gains are due to retrieving analogies. For example, Dugosh and colleagues (2000) found that if people heard lists of ideas as they were brainstorming, their own performance was improved relative to those not hearing lists of ideas, and the more ideas presented the better their performance. These sorts of cues are often fast and cheap to generate. Yet the more arbitrary the cues and the less kind the world, the less efficient this mechanism will be in triggering useful analogical retrievals. Random and arbitrary cues are nonetheless commonly discussed by practitioners, perhaps because of the effectiveness of variable reinforcement schedules for generating superstitious learning (Ferster & Skinner, 1957), or more charitably, because the value of retrieving analogies for creativity and innovation is sufficiently high that any improvement from people's typically low baseline performance is notable. In effect, this work highlights the straightforward claim that if the base rates for analogical retrieval are low, making more retrieval attempts should be more effective than making fewer retrieval attempts. Still, what is most important for current purposes is that this intervention to trigger analogical retrieval occurs at the time of retrieval,

rather than at the time of initial learning, as was commonplace in the problem solving research discussed earlier.

A second area of creativity and innovation research has examined social factors that are predictive of retrieving analogies. As just discussed, providing multiple different cues to memory is one means of accessing information from multiple different areas of memory and improving one's odds given a low overall base rate of analogical retrievals. Another means for making multiple retrieval attempts is to present the same cue to multiple people—and ideally people who have sufficiently different memories such that their retrievals are non-overlapping. For example, one might ask a group of microbiology researchers to explain a puzzling research finding and find that the more diverse the specialty areas of the researchers, the more likely it is that a researcher generates an analogy from one of the related specialty areas that provides an appropriate explanation (cf. Dunbar, 1995).

The idea of tapping diverse groups of individuals to spur creativity extends beyond retrieving analogies, but many of those working with diverse groups (and who are often not, at least at the outset, analogy researchers) tend to end up emphasizing the retrieval and application of analogies in discussing their findings (Burt, 2004; Dunbar, 1995; Hargadon & Sutton, 1997; Paletz & Schunn, in press). Perhaps the most relevant aspect of these findings for our understanding of the individual retrieval of analogies is that there are gains of having access to multiple domains of knowledge. It is possible that individuals who themselves (1) have competence in multiple domains and (2) consider retrieving examples from each of

those domains should be better able to retrieve analogies than those whose competence or retrieval attempts lie solely within one domain.

A third area of creativity research has examined incubation or preparedness effects (Christensen & Schunn, 2005; Moss, Kotovsky & Cagan, 2007; Seifert, Meyer, Davidson, Patalano & Yaniv, 1995). This work, which straddles creativity and problem solving, derives from the observation that people, on rare occasions, spontaneously generate solutions to a problem when they are not actively working on the problem. These are the "a-ha" moments that can strike when walking in the park, getting on a bus, or taking a shower. Here, unlike in problem solving and analogical transfer, the problem or need for creativity comes first, and the solution, which is often an analogy, comes later and retrieves the earlier problem ("oh, my problem can be solved in this same way"). The new aspect for the discussion of analogical retrieval that this work raises is that an unresolved prior example (an open problem or need for creativity) is likely to be retrieved when encountering an analogous example. Thus an additional factor influencing initial encoding of examples that could facilitate later analogical retrieval is whether those examples are finished and settled, or considered still open concerns. Incubation effects suggest that if people have open concerns (or, in the language of social cognition, have concerns that are chronically accessible; Higgins, 1996), this should facilitate analogical retrieval once again simply by leading people to make more retrieval attempts, and retrieval attempts in multiple contexts.

This limited review of research on creativity contributes two useful points to the discussion of analogical retrieval. First, it demonstrates the practical importance

of retrieving analogies and consequently the value of factors that can increase

people's ability to retrieve analogies. Second, it raises the possibility that retrieving

analogies is not only about what is stored. We can also attend to the situation at the

time of retrieval, and consider factors at that point that could spur useful retrievals.

2.3   Acquisition of domain knowledge and retrieving analogies

In discussing problem solving earlier, expertise appeared to be related to

retrieving analogies. Specifically, the prior discussion suggested that expertise

facilitates analogical retrieval because experts encode examples in their domains

using consistent, sophisticated systems of concepts (Forbus, Gentner & Law, 1995).

The question then arises as to how experts learn those sophisticated systems of

concepts that form the heart of their domain knowledge. Drawing analogies likely

plays a key role here (Gentner & Loewenstein, 2002), and there is potentially a key

role for analogical retrieval, as well.

One possibility for why learning sophisticated systems of concepts might

hinge on analogical retrieval is that such concepts must apply to examples that do

not share surface properties. For example, within the domain of negotiation, there

are a relatively small number of kinds of contract structures that enable parties to

form agreements that are beneficial for all involved, and those contract structures

apply to a limitless number of negotiation contexts. A straightforward contract

structure along these lines is called a "tradeoff," and it involves parties conceding on

low priority issues so they can reap concessions from other parties on high priority

issues. Tradeoffs appear in a wide array of negotiation situations. A new graduate

negotiating a job offer might concede and accept a limited benefits package (which

is costly to employers but less pressing to the young and healthy) in return for gaining a high starting bonus to help pay for pressing needs such as a home and car. Or, the US might drop foreign tariffs in return for a country's human rights concessions. Recognizing that varied examples are instances of the same underlying structure—and so forming a relational category (Gentner & Kurtz, 2005)—could involve analogical retrieval to link the varied instances.

To clarify, the use of the term "experts" here is not limited to chess experts, radiological experts, or other experts in advanced domains, but rather applies more generally to all developed domain knowledge. The complex declarative learning central to acquiring domain knowledge is perhaps the major accomplishment of cognitive development as well as the development of expertise in a formal domain (Chi & Ohlsson, 2005; Mason, 2007). The research challenge in studying complex declarative learning is to understand conceptual change or knowledge restructuring. That is, the goal is to understand the learning that results in knowledge that is qualitatively different than what came before and revises one's framework for understanding the domain, rather than simply being the accretion of factual details added to an existing framework (Carey, 1991; Stern, 2005). Although there is still disagreement and a lack of knowledge about the mechanisms underlying conceptual change (Lewandowsky, Kalish & Griffiths, 2000), analogy is on most shortlists (Chi & Ohlsson, 2005; Gentner, Brem, Ferguson, Markman, Levidow, Wolff & Forbus, 1997; Smith, Solomon & Carey, 2005; Vosniadou & Brewer, 1987). The reason is that drawing analogies is a mechanism that can produce generalizations of the underlying structures in examples (Gentner & Wolff,

2000; Hummel & Holyoak, 2003), structures that people might previously have thought to be idiosyncratic rather than being organizing principles or relational categories. For example, novice negotiators may recognize that they formed an agreements specifying that the payments for a television show vary based on the show's eventual ratings, but fail to realize that this is a general type of contract structure—contingent contracts—that can be used in a broad array of negotiation situations (Bazerman & Gillespie, 1999). Because analogies can lead people to recognize generic underlying structure in examples, this means that analogy can produce new representational resources for encoding examples. In brief, analogy plays a fundamental role in the acquisition of domain knowledge because once people are thinking about an analogy, it can lead them to generate sophisticated concepts and systems of concepts (see Gentner & Wolff, 2000).

The question then becomes how people come to draw the analogies that yield new generalizations about domain principles. There are at least two established sources of analogies. One source is progressive alignment (Gentner, Loewenstein & Hung, 2007; Kotovsky & Gentner, 1996): by comparing items that are similar both in their surface properties and underlying principles, people are sensitized to both kinds of commonalities, and become increasingly able to notice other examples sharing the same underlying principles, even if they have different surface properties. For example, children who are asked to compare OXO to oxo are better able to notice a similarity with ABA than children just seeing OXO (and not also oxo) initially (Kotovsky & Gentner, 1996). Progressive alignment uses comparison to improve the quality of people's initial encodings, thereby facilitating

19

later drawing analogies. The difference is that it starts with items with similar surface properties, removing the need for analogical retrieval to initiate the process.

A second basis for people drawing analogies to promote learning domain knowledge is social guidance (Carey, 2004; Gentner & Loewenstein, 2002; Tomasello, 1999). People who are relative experts in a domain can guide learners to draw appropriate analogies and thereby foster their knowledge acquisition. This can be done directly, though formal or informal instruction. It can also be done indirectly, by using words that stand for the important underlying principles in the domain consistently across examples, and thereby guiding learners to compare examples labeled with the same words. For example, it is not always obvious what it means to cooperate, but tracking how the word cooperation is used across situations could guide people to understand this complex kind of social interaction (Keller & Loewenstein, in press). The words make the world kinder for learners, in the sense that they serve to add a salient surface property (Yamauchi & Markman, 2000) to examples with common underlying principles.

What is striking about these two means for using analogies is that they, like the earlier work in problem solving, largely rely on people's encodings of the initial examples. They promote analogical retrieval by encoding examples in useful ways from the start, leveraging people's ability to link items on the basis of surface features. This raises the question of what happens to all the examples people learn that are not encoded in useful ways from the start. These examples might be only minimally useful, mostly lying latent as inert knowledge in memory because they are unlikely to be recalled by an analogy and linked to other examples on the basis

of common underlying principles. If the knowledge experts have is more abstract, interconnected and consistent than novices' knowledge (Chi & Ohlsson, 2005), and is at least in part incommensurable with prior understandings (Smith et al, 2005), then examples learned early may indeed be mostly wasted. They would be mostly wasted unless it is possible for people to revisit earlier examples and re-encode them using later, more sophisticated understandings (cf., de Leeuw & Chi, 2003). Perhaps people can leverage new understandings by integrating them with prior knowledge, and thereby update and learn from that prior knowledge. Perhaps new learning can not only influence what one can do in the future but also one's understanding of the past. It could, if people could retrieve analogies.

3.   Underlying structure and retrieving analogies

One of the most central claims about human memory retrieval (Higham, 2002) is that it is a function of *encoding specificity* (Tulving, 1983; Tulving & Thompson, 1973), or the match between what was initially encoded about some exemplar and the probe that is later used to cue memory retrieval. As a simple example, MBA students who study tables of company financial information are better able to recall that information later if they are presented a similarly formatted empty table rather than an empty table whose rows and columns were transposed (Ryack & Kida, 2006). Encoding specificity has implications for retrieving analogies. As already noted, analogies rest on a match between the underlying structures in two exemplars—the core insight articulated and

elaborated on by the Structure-Mapping theory of analogy (Gentner, 1983; Gentner & Markman, 1997). Thus retrieving analogies from memory must result from people encoding something about the underlying structures in both the stored and probe exemplars, as otherwise there would be no basis for an analogical match. Examining whether, when and how people encode underlying structure in examples is therefore an important concern. Critically though for current purposes, encoding specificity is concerned with a match between what is stored and the cue or memory probe. Although nearly all prior research on analogical retrieval focuses on how stored examples are encoded, encoding specificity provides equal motivation to be concerned with how memory probes are encoded.

3.1 Encoding the underlying structure in examples

There are innumerable ways to encode exemplars, and given that people use their knowledge in multiple ways, people have good reasons for encoding exemplars in multiple ways (Markman & Ross, 2003). For example, for pragmatic reasons people frequently draw attention to different aspects of and different perspectives on exemplars. We might, based on the situation, decide to call a person "the company CFO," "my boss," "an accountant," "an African-American," or "a member of the Executive Committee" and our choice conveys a particular encoding of the situation at hand (Clark, 1996). Such encoding choices should be consequential for memory retrieval, because of encoding specificity.

Young children, novices early in learning a domain, and the lazy or rushed (Gentner & Rattermann, 1991; Hammond, Seifert & Gray, 1991) may encode little or no underlying structure in examples, but otherwise it is reasonable to assume that

22

people encode some underlying structure. For example, 5-6-year-olds tend to generate interpretations of the metaphor "A tape recorder is like a camera," that are focused on surface properties (e.g., they are both black), but 9-10-year-olds and adults generate interpretations that are focused on underlying structure (e.g., they both record something for later; Gentner, 1988). The question then turns to the sources of variation and normal tendencies in how people encode underlying structure. It is already clear from the various interventions that advance problem solving discussed earlier that if emphasizing underlying structure makes a difference for performance, it must be the case that people routinely fail to emphasize the underlying structure in examples.

I suggest that there are at least four sources of variation in encoding underlying structure: ambiguity, context specificity, completeness and weighting. The first is the same ambiguity and contextual variability seen in the CFO/ boss/ etc. example. For example, the act of shaking hands can be interpreted as a greeting (like a wave or a spoken "hello") or as a signal of agreement (like signing a contract). Variation in encoding the underlying structure in examples implies a challenge in choosing which of the multiple possible underlying structures to encode. In most computational models of analogical retrieval (and I will use MAC/FAC, Forbus, Gentner & Law, 1995, and LISA, Hummel & Holyoak, 1997, as running examples because they are arguably the most prominent models), this kind of variation would be instantiated as a difference in which representational units or predicates would be used in a knowledge representation to account for the underlying structure. For example, MAC/FAC takes as input predicate logic representations, composed of

23

units such as GREET(party A, party B) or AGREE(party A, party B). The issue of ambiguity and which underlying structure to encode is unavoidable, and has influences at multiple levels of analysis. Gentner's (1982; Gentner & Boroditsky, 2001; Gentner & Bowerman, 2009) *natural partitions*, *relational relativity* and *typological prevalence* hypotheses map out in large scale reasons for variation and the role of language in guiding people towards some underlying structures as opposed to others. People's own areas of expertise as well as their particular needs and goals in a given situation must also influence what they encode.

The second source of variation in encoding underlying structure is how context specific it is. Underlying structure can be encoded in ways that are specific to the example or in ways that are more generic. For example, one description of a negotiation situation might be: "The Management at Arlington Inc. proposed a 10% pay cut for all hourly workers, which prompted those workers to announce they were considering going on strike." A more generic description of this situation might be: "One party asked for a concession, and the other party refused and made a threat." Computational models of analogical retrieval could instantiate this kind of variation as a difference in which predicates are used in the knowledge representation just as with ambiguity and contextual variability. The difference is that here, more generic descriptions would necessarily draw upon predicates that have different distributions in the larger pool of examples used in the simulation, such that more generic predicates are used in a wider array of contexts, and take a broader range of arguments. If distributed representations were used to encode examples (as in Hummel & Holyoak, 1997), then an additional difference would

likely be that more generic predicates would have fewer components or features. The reason for variation in context specificity is that both specificity and generality are functional. Specificity retains more information and is more tightly linked to the situations of known relevance, while generality provides the potential for powerful and efficient systems of rules. Medin and Ross (1989) outlined reasons why people, and particularly relative novices, are likely to have a bias towards context specific encodings rather than generic encodings. As a result, people likely preserve at least some of the basis for noting multiple underlying structures (i.e., preserve the basis for ambiguity, the first source of variation). However, this means that the underlying structures in examples are likely to be latent rather than manifest (Clement, Mawby & Giles, 1994) in people's representations.

The third source of variation in encoding underlying structure is how complete it is. For example, the use of tradeoffs to form negotiated agreements, as noted earlier, involves exchanging concessions on low priority issues for gains on high priority issues. After reading an example exemplifying a tradeoff, some of our participants (Gentner et al., 2009 Experiment 3) seemed to encode only part of the underlying structure, writing descriptions such as "both parties conceded," "they compromised," "everyone benefitted," or "they were creative." Another form of partial structure is to encode the lower-level events, but fail to consider the overarching pattern (or high-order relational structure) that brings coherence to those events (Clement et al, 1994, study 3; Gentner & Toupin, 1986; Loewenstein & Gentner, 2005). Computational models of analogical retrieval would instantiate this kind of variation as a difference both in how many predicates are used to represent

the underlying structure in the example and in the elaborateness of the system of

relations among those predicates. I am not aware of any general claims about the

completeness of one's encodings of underlying structure in examples. Certainly

one's skill, motivation and opportunity are plausible factors influencing

completeness. The conservative claim, given bounded rationality (Simon, 1947),

would have to be that people tend to encode only part of the underlying structures

in examples.

The fourth source of variation in encoding underlying structure is the weight

or importance placed on that structure relative to other aspects of the example. For

example, greetings may play a very important or a trivial role in a negotiation

example. This could be instantiated in computational models of analogical retrieval

by varying the number of predicates and elaborateness of the system of relations

involving those predicates, as with the discussion of completeness. Encoding

specificity implies that the match between the entire probe example and stored

example is what matters. As a result, the more of the example that is dedicated to

encoding the underlying structure, the more the underlying structure should

determine what matches the example. For example, in MAC/FAC, each component in

a representation has an associated weight that indicates its importance in the initial

matching process (the MAC stage), and the weight of all the components are

normalized such that together they sum to one. It follows that adding more surface

properties to a knowledge representation necessarily means lowering the weight

given to the predicates representing the underlying structure. Or, on the flip side,

simplifying a knowledge representation by removing surface properties necessarily

implies adding weight to the predicates representing the underlying structure. Given the earlier claim that people, particularly relative novices, have a tendency to encode examples in context specific ways, it follows that their encodings maintain surface properties and hence give relatively low weightings to underlying structure.

The picture of people's encodings of underlying structure in examples that emerges from this discussion is one in which relative novices generate context specific encodings. They probably only encode part of the underlying structure, grant it relatively low weight because they maintain ample surface properties, and encode the underlying structure in a way that leaves latent the similarity between the example and other examples of the same underlying structure set in other contexts. Having encoded examples in these ways, it is relatively uncommon for them to retrieve analogies. In contrast, relative experts likely do encode examples using their underlying structures and likely mostly disregard surface properties (Myles-Worsley, Johnston & Simons, 1988). Having encoded examples in this way, it should be relatively more common for experts to retrieve analogies. Finally, the manipulations that seemed to influence analogical retrieval in problem solving and creativity are readily interpretable as ones that influence the encoding of (at least initial) examples in ways that address typical novice shortcomings.

3.2 Using underlying structure in retrieval

Given that underlying structure is at least sometimes encoded, and that encoding specificity implies retrievals should be governed by what is encoded, we must predict that analogical retrieval can and does occur on the basis of underlying structure. The most basic component of structure is role bindings: most parents

want to know if it was John who hurt Mary, or Mary who hurt John. Wharton and colleagues (1994) showed that for sentences almost as simple as these, people were capable of retrieving earlier sentences with not just similar content (Jack, Marie, hit), but also matching role bindings (that Jack was the one doing the hitting, and it was Marie who got hit).

A second important concern is whether people can retrieve prior examples based on many relations, not just one. An elegant demonstration by Clement, Mawby and Giles (1994) demonstrated that people can, and that generic encodings of underlying structure matter. They showed that people were less likely to retrieve an analogy if the examples were written with context specific verbs (e.g., verbs specific to writing, such as edit… marked… typed in; verbs typically encode the relations that make up underlying structure) than if the examples were written with generic verbs (e.g., verbs not specific to writing that can apply to many contexts, such as fixed… replaced). Using generic verbs served to facilitate people's analogical retrievals.

A final concern is whether people can retrieve analogical examples sufficiently well to make use of them. Ross and Kennedy (1990) showed that people who first linked two examples with a common underlying structure were later able to use these early examples to help solve a subsequent problem. Perhaps the most evidence in support of the use of underlying structure in retrieval are the studies showing that people who draw comparisons between examples and derive generalizations of their commonalities, relative to those who see just one initial example or two examples separately, are later more likely to solve an analogous problem (e.g., Catrambone & Holyoak, 1989; Gick & Holyoak, 1983; Loewenstein,

Thompson & Gentner, 1999, 2003). The explanation for these effects rests on people retrieving either the earlier examples or a generalization derived from those examples upon encountering the subsequent analogous problem—it rests on analogical retrieval. The question is, how does analogical retrieval work.

The approach taken by the main models of analogical retrieval generally starts from content matches and moves to structural matches that take role bindings into consideration. MAC/FAC (Forbus et al., 1995) uses two distinct phases: one in which all content in a probe example is matched against all content of examples in memory without consideration for structure, and then the most similar items from memory are examined in a second phase that takes the entire structure of the examples into consideration. LISA (Hummel & Holyoak, 1997) uses less distinct phases, matching subsets of the content in examples in separate waves (in effect using content and a little information about structure) and then subsequently adds in a fuller accounting of structural consistency. A further, subtler commonality across the models is that they normalize the weights of the items within a single knowledge representation—forcing each example to have a total weight of one, rather than having each predicate within a representation have a weight of one—to avoid examples with large numbers of predicates from being constantly retrieved (Forbus et al., 1995; Hummel & Holyoak, 1997). As a result of this normalizing, not just the number of predicates articulating the underlying structure but their proportion relative to the rest of the example should therefore influence the rate of retrieving analogies. Accordingly, the most basic explanation that models of analogical retrieval give to explain why people retrieve analogies is that examples

29

are encoded in generic rather than context specific form, with complete underlying structures, and with relatively few surface properties so their underlying structures gain sufficient weight.

Why generic encodings facilitate analogical retrieval is complex. One possibility is that generically encoded examples better retrieve other generically encoded examples than contextually encoded examples retrieve other contextually encoded examples (assuming that the examples are from different contexts). This was the contrast Clement and colleagues (1994) examined: examples from different contexts were both given generic verbs, and this facilitated analogical retrieval relative to when both examples were given context-specific verbs. Of interest for understanding the mechanism, Clement and colleagues found comparably strong results even if the generic verbs in the two examples were not the same, but merely synonymous.

It could be that people's encodings were identical even if the verbs were only synonymous, but this issue appears in more challenging forms, so it is still worth considering that two different but generic encodings are readily matched. In the work on schemas in problem solving, people derive schemas initially, and then later attempt to solve a problem, set in a specific context. Having generated the schema earlier facilitates problem solving, presumably because of the retrieval and application of the schema. It is possible that the schema is retrieved because having generated a schema makes people more likely to encode the later problem using generic rather than context specific representations of the underlying structure. To some extent this begs the question of how people know to use a more generic

30

encoding of the seemingly specific problem, but perhaps one could argue that the encoding vocabulary has changed as a result of learning the schema. Alternatively, or in addition, it could be that generic encodings not only match other generic encodings, but can also match specific encodings. For example, using the materials of Clement and colleagues (1994), perhaps examples with generic verbs, like fixed and replaced, would retrieve prior analogous examples with context specific verbs, like edit and typed in. This would be an interesting area for future research. Some kind of predicate decomposition (Gentner & Wolff, 2000), minimal ascension to a superordinate category (Falkenhainer, 1990), or distributed representation scheme (as used by the LISA model, Hummel & Holyoak, 1997) seems necessary to allow a generic encoding to match a specific encoding.

The weighting of aspects of a representation also has attendant complexities. Specifically, it suggests that there could be a difference between schemas or generalizations on the one hand and examples on the other. This is because examples will necessarily include surface properties. If the example is sufficiently large and complex, it could also include extraneous underlying structure distinct from the main thrust of the example. If normalizing the total weight of an example to one is an appropriate assumption, then this suggests that removing surface properties and tangential underlying structure will result in increasing the weight on the main underlying structure in the example. This, in turn, will increase the likelihood of the example being retrieved as an analogy. If this is so, the more that examples are turned into schemas or generalizations that encode primarily underlying structure and little in the way of surface properties, the more they

should be able to be retrieved as analogies, and serve as general purpose rules (Gentner & Medina, 1998). Alternatively, it could be that the normalizing assumption needs to be made more complex, such that people can maintain the surface properties in examples but lessen their weight (or learn their irrelevance; Ross & Kennedy, 1990). This would allow examples to function like schemas without forcing them to lose all context specificity.

A more general complexity to note is that analogical retrieval must be able to succeed on the basis of matches of partial underlying structures, rather than entire systems of underlying structure (Johnson & Seifert, 1992; Kurtz & Loewenstein, 2007). For example, the saying "don't count your chickens before they are hatched" might come to mind upon hearing a story about a man making elaborate preparations to spend an inheritance of a still-living relative, even before finding out that the relative did not include the man in her will (cf., Johnson & Seifert, 1992). More generally, analogical problem solving hinges on retrieving prior examples that serve as bases from which people can infer solutions for their current problems. Accordingly, the match between a prior example and a current problem must be based on only partial underlying structures, because the current problem has no solution to include in the matching process. The prior example and the current problem must match on the basis of the description of the problem situation, even if it is the solution that is the purpose of having retrieved the analogy. Schemas might be useful for providing generic descriptions of solutions, but unless people run through a list of all possible solutions they know, this effect can only be important for applying the schema, not for retrieving it. Although analogical retrieval is

relatively rare, it is not because underlying structures need to match completely, just substantively.

The final point to highlight about analogical retrieval is that there is nothing in encoding specificity and in the matching processes in computational models of analogical retrieval that makes them directional. There is a fundamental asymmetry in memory retrieval in the sense that there is one probe and a vast number of items stored in memory, but this is a distinct issue from the process of evaluating any given match between the probe and a given example in memory. Even though it is the match between the probe and the stored, initial example that should matter, nearly all the research on analogical retrieval has examined changes to the encoding of the initial examples. However, all the main concerns (generic encodings, complete underlying structures, minimal surface properties, emphasized or weighted underlying structure) could just as readily be applied to the probe example. As Dedre Gentner put it when having the initial insight, maybe transfer can go backwards. Maybe the same factors that allow comparison and other interventions to foster forward transfer to new problems can also foster backwards retrieval of previously learned analogous examples.

It is possible that transfer does not go backwards, and that prior research has been correct to focus on the encoding of initial examples. Perhaps more generic encodings of early examples change the encoding vocabulary used to encode later examples, and this is why the initial examples are later recalled. A more generic encoding of a probe example presumably would not retroactively change how the content of previously learned examples are encoded, and hence might not be

effective. Still, if encoding specificity guides retrieval, then it is possible that altering

the encodings of probe items to emphasize their underlying structure could

influence the rate of retrieving analogous examples.

4.   Facilitating the retrieval of analogies at retrieval time

My colleagues and I (Gentner, Loewenstein & Thompson, 2004; Gentner et al,

2009; Kurtz & Loewenstein, 2007) set out to examine whether changing people's

encoding of an example could influence their ability to retrieve an analogy already

stored in memory. We chose comparison as a means for changing people's

encodings of an example, as it is well tested and robustly fosters complete and

generic encodings that deemphasize surface properties. Also, we know that people

encounter comparisons in the course of their work, their education, and even

through art and literature; it is possible that those comparisons not only trigger

ideas for the future, but also trigger reflections on prior analogous examples. In our

studies, rather than have people compare two examples then test whether they

were able to retrieve that learning in response to some subsequent example or

problem, we instead asked people to compare two examples and tested whether this

facilitated retrieving a previously learned analogous example. We found that it can.

4.1 The "Own Memory" studies: Retrieving analogies from autobiographical memory

Some research starts in the lab, but this research started by examining

management consultants from a major firm as they were learning negotiation skills

through a professional training seminar. This is a compelling sample and setting for several reasons. We were interested in people's ability to retrieve analogies, and this is a task that management consultants are paid to do. They are general problem solvers, often working in multiple industries, with different kinds of companies, in different parts of the country and even different countries. As a result, they have a professional need to retrieve analogies as they confront new problems. Still, the consultants we worked with were engaged in a training seminar, and so learning about a domain, negotiation, about which they were familiar but not expert. Accordingly, we could still examine them learning and retrieving examples with underlying structures that they had not already mastered. Also, because they were engaged in a training seminar they found important and useful to their careers, and training alongside their peers at their organization, they were highly motivated to learn, to impress their peers, and to avoid embarrassing themselves. Finally, working with this group, we could assess their retrievals of examples learned days, months, even years earlier because they had familiarity with the domain under study. That is, we could examine people's analogical retrievals from their own memories as they struggled to learn something important to them. The limitation, of course, is that we do not know how the stored examples were initially encoded. But this approach does allow us to examine a natural situation of significant importance: given scattered learning over a period of years, what likelihood is there for retrieving an analogy to a current example.

The domain under study was negotiation, which, as was mentioned earlier, is a compelling domain in which to study analogical retrieval because the same

underlying structures recur in widely different contexts. It is also a topic about which our participants were highly motivated to learn. We examined people's performance retrieving tradeoffs in one study, but in most studies, we examined people's performance retrieving contingent contracts, or agreements whose terms vary according to the outcome of a future event. There are many standardized subtypes, such as bets, options contracts, and pay-per-performance contracts, as well as countless idiosyncratic examples. It is a class of negotiated agreement structures with known properties that can be used to motivate performance, manage risk, overcome biased perceptions, bridge timing differences, and test for deception (Bazerman & Gillespie, 1999). In prior research in the negotiation context using contingent contracts and tradeoffs with undergraduates, MBA students and executives as participants, my colleagues and I have found that encoding training examples by drawing comparisons between them facilitates forward knowledge transfer (Gentner, Loewenstein & Thompson, 2003; Loewenstein, Thompson & Gentner, 1999; 2003; Thompson, Gentner & Loewenstein, 2000). In most of these studies, we showed that drawing comparisons facilitated later using the agreement structure from the examples in a face-to-face negotiation.

In the "own memory" studies, we were focused on whether drawing comparisons facilitated retrieving analogous examples from autobiographical memory. In our first study, we randomly assigned consultants to analyze two examples of the contingent contract structure either one example at a time, or both examples together with an instruction to compare. After writing about the examples we provided, the consultants retrieved matching examples. About 10% did not

36

retrieve any example, which, given the sample and the notes many of those 10% wrote, mostly indicates that they could not retrieve an example they felt had a sufficiently close match to the examples we gave them. Analogical retrieval is indeed difficult. Still, the comparison manipulation had a clear effect. In the probe comparison group, 55% retrieved a complete analogy match, whereas 35% in the separate case analysis group did so. In studies replicating and extending the effect with less experienced samples, we found similarly substantial differences: 74% versus 29% for masters of accounting students (with about 1 year of work experience), and 50% versus 16% for MBA students (with about 5 years of work experience; these last percentages are collapsed across two rounds of retrievals, one for contingent contract examples and one for tradeoff examples). Drawing comparisons between memory probes at the time of retrieval facilitates analogical retrieval from one's own memory.

The data further support linking analogical retrieval to the quality of participants' encodings of the memory probes, just as the quality of participants' encodings of training examples has been shown to predict forward transfer (e.g., Gick & Holyoak, 1983; Loewenstein, Thompson & Gentner, 1999). From participants' written descriptions of the examples, we derived a measure of the completeness of their understandings of the underlying structure in the examples. (It would be interesting in further research to code how generic people's descriptions are and how many surface properties they list.) We found that on average, the understandings from the probe comparison group were rated as more complete than the understandings from the separate case analysis group, and that

completeness was correlated with analogical retrieval. Thus across over 250 participants with three different levels of domain experience, we have clear support that drawing a comparison between examples spurs more complete descriptions of the underlying structures, and in turn more analogical retrievals (Gentner et al., 2009). Fostering more complete understandings of the underlying structure in probe examples can facilitate retrieving analogies from one's own memory.

4.2 The controlled memory set studies

A second set of studies examined analogical retrieval from a controlled memory set, rather than from people's autobiographical memories. This allowed us to examine whether the analogies retrieved from people's own memories were dependent on any special sort of initial encoding. It is possible that the retrieval advantage found in the "own memory" studies only holds for stored examples that people encoded in an unusually effective manner (i.e., with complete, generic underlying structures and few surface properties). The simplest controlled memory set study involved presenting undergraduate participants with a set of seven examples of negotiations, inserting a brief, filled delay, then having them be in either the probe comparison or probe separate cases group (Gentner et al., 2009). Rates of analogical retrieval were considerably lower than in the own memory studies, but still there was an advantage of probe comparison (27% analogical retrieval) over probe separate case analysis (6%). This suggests that even routinely encoded examples can be later retrieved on the basis of underlying structure, provided people have encoded the probe example effectively.

The task in the studies described thus far was simply to recall a prior matching example, but as the earlier discussion indicated, a primary dependent measure used to study analogical retrieval is problem solving transfer. For this reason, Ken Kurtz and I (Kurtz & Loewenstein, 2007) studied retrieval-time encoding effects during problem solving. Specifically, we used Gick and Holyoak's (1980, 1983) classic convergence materials for studying analogical problem solving: Duncker's radiation problem whose solution requires the use of multiple, low-power rays arranged to converge on and eliminate an inoperable tumor, and an analogous story about a general who arranges to have multiple small groups of troops converge on and capture a fortress. The difference between our studies and prior research was that, instead of focusing on how participants encoded the fortress story initially, we focused on how participants encoded the tumor problem. Participants read the fortress story, and then were given one of a variety of tasks. The critical group was asked to compare the tumor problem with a second, analogous problem (Gick and Holyoak's Red Adair story, turned into a problem whose solution hinges on how to extinguish a fire using many small hoses simultaneously). Then participants were asked to solve the problems. If the comparison led participants to generate a more complete or more generic encoding of the underlying structure of the problem, they should be more likely than other participants not engaging in a comparison to retrieve the earlier fortress story and generate convergence solutions. In support of this logic, we found that participants who had compared problems were more likely to generate convergence solutions to the tumor problem than participants who only received the tumor problem (54%

versus 15%; Experiment 1) or who worked on the two problems separately (38% versus 15%; Experiment 2). As an extension, we replicated this second study, contrasting the problem comparison and separate problems conditions, using a roughly 20 minute, filled delay between reading the initial fortress story and later attempting to solve the problems (Gentner et al., 2009). We found a similar pattern of results, albeit with a lower overall rate of solutions. Those who compared the two problems (31%) were more likely to solve the tumor problem than those who read the problem separately (5%). These studies are consistent with the claim that an improved encoding of problem descriptions, not solution descriptions, can facilitate analogical retrieval.

Unlike for studies of forward transfer, when people are comparing stories in support of later problem solving, in these last studies the comparison was between the problems themselves, which meant we needed an additional kind of control. We needed to test whether comparing two problems, in itself, was the basis of the problem solving advantage, rather than it being an advantage because it facilitated retrieving the earlier analogous story. Accordingly, we contrasted two groups that both compared the two problems, but one group did and one group did not first read the original fortress story (Kurtz & Loewenstein, 2007). The evidence points to retrieving the earlier story as the key advantage, as those receiving the fortress story analogy (51%) were more likely to solve the tumor problem than those not receiving it (34%). Further evidence in support of analogical retrieval comes from a post-task questionnaire. Those who had read the fortress story initially were asked how they attempted to solve the problems. Of those generating a convergence

solution to the tumor problem, 53% said they retrieved an earlier story to use as an analogy, whereas only 15% of those who did not generate a convergence solution reported retrieving an analogy. Taken together, these problem solving studies suggest that people can spontaneously retrieve and use a prior, single example on the basis of underlying structure, provided they have an effective encoding of the problem they are trying to solve.

4.3 MAC/FAC simulation modeling

Simulation studies are useful for examining cognitive processes, and we conducted a series of four simulation studies to help examine the plausibility of obtaining different rates of analogical retrievals based on how examples are encoded (Gentner et al., 2009). The goal was to examine whether a consistent retrieval process could account for both the previously established forward transfer effects based on using comparison to encode training examples stored in memory and the new analogical retrieval effects based on using comparison to encode examples used as memory probes.

It was not feasible to simulate the "own memory" studies directly; generating vast numbers of knowledge representations in a new domain is non-trivial. Instead, we took the approach of addressing the theoretical issues using arguably the largest and most widely used set of knowledge representations used in analogical retrieval studies, the "Karla the Hawk" materials (Gentner, Rattermann & Forbus, 1993). We added to this story set by generating additional analogs (there were already two analogs per story set; we added a third for each set), and generating schemas (one per story set, from two of the analogs). Then we examined the retrievals generated

41

if a story and if a schema were used as probe items (to test the new analogical

retrieval effects), or if a schema was stored in memory (to test forward transfer

effects). Consistent with the experimental studies, we showed that schemas used as

probe items tended to retrieve analogous examples, schemas used as memory items

tended to be retrieved by examples used as probes, and examples as probe items

tended to retrieve examples with similar surface properties rather than analogous

examples.

The interesting part of simulation studies though is why they generate the

pattern of results they do. In these simulations, schemas had complete

representations of the underlying structure in examples, and few surface properties.

Examples also had complete representations of the underlying structure, and many

surface properties. Thus in these simulations, surface feature matches and low

weights on underlying structure are what led the examples to be relatively

ineffective for analogical retrieval. The absence of surface feature matches and the

high weights on underlying structure are what led the schemas to be effective for

analogical retrieval, both when serving as probes retrieving analogous examples,

and when in memory being retrieved by analogous examples. We did not simulate,

but it would be straightforward to show, that MAC/FAC would generate a similar

pattern if we had manipulated the completeness of the underlying structures rather

than their weights. It is less straightforward to examine the role of generic versus

context specific representations of the underlying structure in examples; this would

be a worthwhile endeavor for future modeling work to consider. Even so, just from

these simulations, it is notable that the same mechanisms that can account for

schema abstraction and forward transfer can be parsimoniously extended to account for probe encoding and analogical retrieval. Forward transfer need not be entirely due to learning to encode examples more effectively or due to encoding solutions effectively.

Taken together, the theoretical, experimental and simulation work my colleagues and I have done combine to establish a new avenue for research. It is possible to facilitate analogical retrieval without changing how people initially encode examples. This is not to say that the quality of initial encodings is irrelevant; that is just an already established point. Instead, it was an open question as to whether people needed to have high-quality encodings of stored examples to be able to retrieve those examples later on the basis of underlying structure. This does not appear necessary. Effectively encoding current examples supports analogical retrieval to examples encoded in mundane ways.

5.   Implications

The most straightforward implication of the findings on comparing probe items facilitating analogical retrieval is that inert knowledge can be revived. Examples learned prior to completely or generically understanding underlying structures and prior to understanding which surface properties are unimportant can nonetheless be retrieved later on the basis of an analogy to a well encoded probe. To elaborate on this central implication, it is useful to return to the earlier discussions of problem solving, creativity and the acquisition of domain knowledge.

43

5.1 Implications for problem solving and creativity

The studies just discussed showed effects of comparing examples and problems on analogical retrieval and problem solving. One striking possibility raised, particularly by the Kurtz and Loewenstein (2007) studies, is that sometimes having two problems is better than having one problem. Still, in these studies participants were provided with problems or examples to compare. It is possible for people to generate or notice effective comparisons themselves, but it seems far more likely that most of the time, people will have a single current situation or problem to address. Consequently, the key point for problem solving and creativity may well turn out to be that these studies help to establish a pathway—altering encodings of probe examples can increase analogical retrievals—that further work shows how to exploit using interventions apart from comparison that individuals can more readily implement themselves when attempting to solve problems or generate creative outcomes.

I have discussed multiple interventions that can influence the completeness of, generic articulation of, and emphasis on underlying structure (cf., Burstein, Collins & Baker, 1991). Self-explanations, fading out of surface properties and other methods for de-emphasizing surface properties, words that articulate underlying structure, and oral presentation of examples might also all facilitate analogical retrieval by influencing the encoding of a current problem or situation. For example, writing out one's problem and then systematically replacing the context specific verbs with more generic verbs, following Clement and colleagues' (1994) work, could enhance analogical retrieval in the service of deriving new solutions or

creative outcomes. Studying these various interventions and assessing the different kinds of variation in encoding they induce would provide the basis for establishing a flexible array of interventions that people could use, on their own, to facilitate analogical retrieval and thereby advance problem solving and creativity.

Within research on creativity and innovation, there is a tradition of focusing on how people are encoding the task or problem at hand (e.g., Csikszentmihalyi & Getzels, 1971). But it is mainly been more recently that encodings of one's problem have been clearly linked to analogical retrieval. Markman and colleagues (2009) discuss several means for fostering what I have been calling complete, generic encodings of problems' underlying structure. Ward (2009) emphasizes using synonyms and category taxonomies to generate multiple alternative encodings of a problem. This should enable people to capitalize on the ambiguity of the problem's underlying structure and retrieve a broader array of analogies. Analogical retrievals in the service of creativity and innovation need not be left to chance, but can be directly and systematically fostered.

A further possibility, over and above the prior mentioned interventions for influencing people's encodings, is to consider the potential effects of construal level (Trope & Liberman, 2003). Construal level theory suggests that, roughly, if people perceive events to be distant, they will encode them in more abstract ways, whereas if people perceive events to be proximal, they will encode them in more detailed ways. Construal level influences similarity judgments (Day & Bartels, 2008). Consequently, it is feasible that construal level could influence probe encodings and hence analogical retrieval. It is possible that distant construals would yield more

analogical retrievals than proximal construals, if distant construals place greater weight on abstract, underlying structure or reduce attention to surface properties. The caveat is that distance could lead to over-simplification.

All of these are suggestions about how to use the pathway established by the new research showing that changing the encoding of a current example can change the likelihood of retrieving analogies from memory. The studies themselves presented people with appropriate comparisons as a means to usefully alter their encodings of current examples, but there are multiple alternative interventions apart from comparison that could offer gains and be easier for individuals to implement on their own.

5.2   Implications for the acquisition of domain knowledge

An exciting phenomenon that arises in teaching is when a student, upon learning some principle, has an immediate reminding: "that explains a puzzling event I never understood" or "that happened to me once!" The studies just discussed provide an explanation for this phenomenon. Students are being guided to generate effective representations of a current example that then triggers an analogical retrieval. More generally, the phenomenon points to a new occasion and pathway for integrating domain knowledge.

If generating more effective encodings of current examples makes those examples more useful bases for retrieving analogous examples from memory, then any occasion of insight into a current example can be a trigger to useful reflection. Specifically, a useful comparison generated by an instructor or on one's own, an act of self-explanation, or the use of any of the other interventions previously discussed

could yield a more complete, generic encoding emphasizing underlying structure. This could prompt a reminding to an analogous example that was stored in an ordinary fashion. This in turn could prompt the re-encoding of the prior example using the more complete and generic encoding of the underlying structure. The result could be the beginnings or the further refinement of a sophisticated category for understanding the domain. Although this account is tentative, the studies described earlier have already provided evidence for all but the last of the steps in this chain. This chain could be the basis for ongoing, routine bootstrapping of current insights to revisit and reorganize, rather than leave unused, prior learning. It is a means for knowledge integration.

In emphasizing knowledge integration and a specific chain of events for integrating new insights with old knowledge, I am finessing a debate in the conceptual change literature (e.g., diSessa & Sherin, 1998; Smith, 2007) about what is changing and how much changes in conceptual change. The sort of knowledge integration through analogical retrieval and re-encoding of old knowledge I am discussing could yield minor updates to one's domain knowledge. It could produce changes to domain knowledge slowly, leaving domain knowledge fragmented, as knowledge from various contexts within the domain need not be integrated at the same time. It is also possible that the chain of events I have described could generate differences in degrees of understanding that, after sufficient numbers of passes through the cycle occur, generate qualitative differences in the kind of knowledge people use to encode domain examples. A gradual changing over of knowledge mainly focused on surface properties and context-specific encodings of

47

underlying structure could shift to knowledge mainly focused on generic encodings of underlying structure with qualitatively different possibilities. Finally, it is also possible that the knowledge integration process I have outlined could be a means by which a striking or radical insight of the sort Carey (2004) discusses in the domain of number could have a broad and systemic influence by triggering the analogical retrieval and re-encoding of a whole range of old knowledge. This could be coupled with a broader effort of deliberate self-explanation and reflection-driven conceptual change (e.g., de Leeuw & Chi, 2003). My point is that the mechanism my colleagues and I have described, whereby an effective encoding can retrieve analogous examples that were encoded in less effective ways, provides a means for new understandings to change and integrate prior knowledge, thereby extending the reach and impact of those new understandings.

This discussion is aimed at highlighting a possibility for how conservative learning can be later revisited once people develop more sophisticated domain understandings. It is worth considering how thoroughly the old examples are re-encoded in light of the new, more complete, more generic understandings. The reason is that new insights can be wrong, still later learning can suggest alternative generic encodings, and more generally it is plausible for there to be multiple waves of conceptual change. There are many reasons conceptual change research has tended to focus on basic mathematical and scientific understandings. Yet the theories presumably apply just as much to still changing areas of science, social science and the humanities. These are areas with less stability in the systems of generic underlying structures in use by experts. This line of thinking gives new

reason to consider whether, for example, drawing comparisons alters people's encodings of the examples themselves, generates schemas or abstract principles that are distinct from the examples themselves, both, or something else. The effects may be highly similar for the local act of solving a current problem, but they may quite distinct at the level of acquiring and re-acquiring domain expertise, particularly for still-changing domains. The recent research my colleagues and I have done suggests people do not have to learn it right the first time, but it is not yet clear how many times people can re-learn it.

6. Conclusion

Memory provides an ocean of examples from which people draw to act intelligently. Usually people reach close at hand for an example, and for many mundane tasks this yields satisfactory outcomes. But for more challenging problems, for which expertise matters, people often need to cast their lines farther to retrieve a useful analogy. Prior research has emphasized efforts at bringing those other examples closer at hand. The new research presented here suggests it is also possible to better bait one's hook such that a longer cast catches a keeper.

References

Ahn, W. K., Brewer, W., & Mooney, R. (1992). Schema acquisition from a single example. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 391-412.

Amabile, T. M. 1996. *Creativity in Context.* Boulder, CO: Westview Press.

Bassok, M., & Holyoak, K. J. (1989). Interdomain transfer between isomorphic topics in algebra and physics. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 153-166.

Bazerman, M. H., & Gillespie, J. J. (1999). Betting on the future: The virtues of contingent contracts. *Harvard Business Review, 77*(5), 155-160.

Benyus, J. M. (1997). *Biomimicry: Innovation inspired by nature*. New York: William Morrow.

Blessing, S. B., & Ross, B. H. (1996). Content effects in problem categorization and problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22(3),* 792-810.

Brown, A. L. (1989). Analogical learning and transfer: What develops? In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 369-412). New York: Cambridge University Press.

Brown, A. L. (1990). Domain-specific principles affect learning and transfer in children. *Cognitive Science, 14*, 107-133.

Burstein, M. H., Collins, A., & Baker, M. (1991). Plausible generalization: Extending a

model of human plausible reasoning. *The Journal of the Learning Sciences,*

1(3&4), 319-359.

Burt, R. S. (2004). Structural holes and good ideas. *American Journal of Sociology,*

110(2), 349-399.

Campbell, D. T. (1960). Blind variation and selective retention in creative thought as

in other knowledge processes. *Psychological Review, 67*(6), 380-400.

Carey, S. (1991). Knowledge acquisition: Enrichment or conceptual change?. In S.

Carey & R. Gelman (Eds.), *The epigenesis of mind: Essays on biology and*

cognition (pp. 257–291). Hillsdale, NJ: Erlbaum.

Carey, S. (2004). Bootstrapping and the origin of concepts. *Daedalus*, 59–68.

Catrambone, R., & Holyoak, K. J. (1989). Overcoming contextual limitations on

problem-solving transfer. *Journal of Experimental Psychology: Learning,*

Memory and Cognition, 15(6), 1147-1156.

Chi, M. T. H., Feltovitch, P. J., Glaser, R. (1981). Categorization and representation of

physics problems by experts and novices. *Cognitive Science, 5*, 121-152.

Chi, M. T. H., & Ohlsson, S. (2005). Complex declarative learning. In K. J. Holyoak, & R.

G. Morrison (Eds.), *The Cambridge Handbook of Thinking and Reasoning* (pp.

371-399). New York: Cambridge University Press.

Christensen, B. T., & Schunn, C. D. (2005). Spontaneous access and analogical

incubation effects. *Creativity Research Journal*, 17(2), 207-220.

Christensen, B. T., & Schunn, C. D. (2009). "Putting blinkers on a blind man":

Providing cognitive support for for creative processes with environmental

cues. In A. B. Markman & K. L. Wood (Eds.), *Tools for innovation: The science behind the practical methods that drive new ideas* (pp. 48-74). New York: Oxford University Press.

Clark, H. H. (1996). *Using language*. New York: Cambridge University Press.

Clement, C. A., Mawby, R., & Giles, D. E. (1994). The effects of manifest relational similarity on analog retrieval. *Journal of Memory and Language, 33,* 396-420.

Csikszentmihalyi, M., & Getzels, J. W. (1971). Discovery-oriented behavior and the originality of creative products: A study with artists. *Journal of Personality and Social Psychology, 19*(1), 47-52.

Day, S. B., & Bartels, D. M. (2008). Representation over time: The effects of temporal distance on similarity. *Cognition, 106*, 1504-1513.

de Leeuw, N. & Chi, M.T.H. (2003). The role of self-explanation in conceptual change learning. In G. Sinatra & P. Pintrich (Eds.), *Intentional Conceptual Change* (pp. 55-78). Mahwah, NJ: Erlbaum.

DeLoache, J. S. (1995). Early symbol understanding and use. In D. Medin (Ed.). *The psychology of learning and motivation (Vol. 32).* New York: Academic Press.

diSessa, A., & Sherin, B. (1998).What changes in conceptual change? *International Journal of Science Education*, 20, 1155–1191.

Dixon, J. A., & Bangert, A. S. (2004). On the spontaneous discovery of a mathematical relation during problem solving. *Cognitive Science, 28*, 433-449.

Dugosh, K. L., Paulus, P. B., Roland, E. J., & Yang, H.-C. (2000). Cognitive stimulation in brainstorming. Journal of Personality and Social Psychology, 79(5), 722-735.

Dunbar, K. (1995). How scientists really reason: Scientific reasoning in real-world

    laboratories. In R.J. Sternberg, & J. Davidson (Eds.). *Mechanisms of insight*.

    Cambridge MA: MIT press. pp 365-395.

Dunbar, K., & Blanchette, I. (2001). The in vivo / in vitro approach to cognition: The

    case of analogy. *TRENDS in Cognitive Sciences, 5*(8), 334-339.

Falkenhainer, B. (1990). A unified approach to explanation and theory formation. In

    J. Shrager & P. Langley (Eds.), *Computational models of scientific discovery and*

    theory formation (pp. 157-196). Los Altos, CA: Morgan Kaufmann.

Ferster, C. B., & Skinner, B. F. 1957. *Schedules of reinforcement*: American

    Psychological Association.

Finke, R. A., Ward, T. B., & Smith, S. M. (1992). *Creative Cognition: Theory, Research,*

    and Application. Cambridge, MA: MIT Press.

Forbus, K. D., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based

    retrieval. *Cognitive Science, 19*(2), 141-205.

Gentner, D. (1982). Why nouns are learned before verbs: Linguistic relativity versus

    natural partitioning. In S. A. Kuczaj (Ed.), *Language development: Vol. 2.*

    Language, thought and culture (pp. 301-334). Hillsdale, NJ: Lawrence

    Erlbaum Associates.

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. (1989). Mechanisms of analogical learning. In S. Vosniadou, & A. Ortony, (Eds.), *Similarity and Analogical Reasoning* (pp. 199–241). New York: Cambridge University Press.

Gentner, D., & Boroditsky, L. (2001). Individuation, relativity and early word learning. In M. Bowerman & S. Levinson (Eds.), *Language acquisition and conceptual development* (pp. 215-256). Cambridge, UK: Cambridge University Press.

Gentner, D., & Bowerman, M. (2009). Why some spatial semantic categories are harder to learn than others: The Typological Prevalence hypothesis. In J. Guo, E. Lieven, S. Ervin-Tripp, N. Budwig, S. Özçaliskan, & K. Nakamura (Eds.), *Cross-linguistic approaches to the psychology of language: Research in the tradition of Dan Isaac Slobin* (pp. 465-480). New York: Lawrence Erlbaum Associates.

Gentner, D., Brem, S., Ferguson, R. W., Markman, A. B., Levidow, B. B., Wolff, P., & 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., & Kurtz, K. (2005). Relational categories. In W. K. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman & P. W. Wolff (Eds.), *Categorization inside and outside the lab.* (pp. 151-175). Washington, DC: APA.

Gentner, D., & Loewenstein, J. (2002). Relational language and relational thought. In E. Amsel & J. Byrnes (Eds.), *Language, Literacy, and Cognitive Development: The Development and Consequences of Symbolic Communication* (pp. 87-120). Mahwah, NJ: Erlbaum.

Gentner, D., Loewenstein, J., & Hung, B. (2007). Comparison facilitates children's

    learning of names for parts. *Journal of Cognition and Development*, 8(3), 285-

    307.

Gentner, D., Loewenstein, J., & Thompson L. (2003). Learning and transfer: A general

    role for analogical encoding. *Journal of Educational Psychology, 95*(2) 393–

    408.

Gentner, D., Loewenstein, J., & Thompson, L. (2004). Analogical encoding:

    Facilitating knowledge transfer and integration. In K. Forbus, D. Gentner, and

    T. Regier (Eds.), *Proceedings of the Twenty-Sixth Annual Conference of the*

    Cognitive Science Society (pp. 452-457). Chicago, Il: Cognitive Science Society.

Gentner, D., Loewenstein, J., Thompson, L., & Forbus, K. D. (2009). Reviving inert

    knowledge: Analogical abstraction supports relational retrieval of past

    events. *Cognitive Science, 33*, 1343-1382.

Gentner, D., & Markman, A. B. (1997). Structure mapping in analogy and similarity.

    American Psychologist, 52, 45-56.

Gentner, D., & Medina, J. (1998). Similarity and the development of rules. *Cognition,*

    65, 263-297.

Gentner, D., & Rattermann, M. J. (1991). Language and the career of similarity. In S.

    A. Gelman & J. P. Brynes (Eds.), *Perspective on thought and language:*

    Interrelations in development (pp. 225–277). New York: Cambridge

    University Press.

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

Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, 10, 277-300.

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

Gick, M. L., & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology, 12*, 306-355.

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

Goldstone, R. L., & Son, J. Y. (2005). The transfer of scientific principles using concrete and idealized simulations. The *Journal of the Learning Sciences*, *14*(1), 69-110.

Goldstone, R. L., & Wilensky, U. (2008). Promoting transfer through complex systems principles. *Journal of the Learning Sciences, 26*(1), 465-516.

Hammond, K. J., Siefert, C. M ., & Gray, K. C. (1991) . Functionality in analogical transfer: A hard match is good to find. *The Journal of the Learning Sciences*, 1, 11-152.

Hargadon, A. B., & Sutton, R. I. (1997). Technology brokering and innovation in a product development firm. *Administrative Science Quarterly*, 42, 716-749.

Hesse, M. B. (1966) . *Models and analogies in science*. Notre Dame, IN: University of

    Notre Dame Press.

Higgins, E. T. (1996). Knowledge activation: Accessibility, applicability, and salience.

    In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of*

    basic principles (pp. 133–168). New York: Guilford Press.

Higham, P. A. (2002). Strong cues are not necessarily weak: Thompson and Tulving

    (1970) and the encoding specificity principle revisited. *Memory & Cognition,*

    30(1), 67-80.

Hofstadter, D., & the Fluid Analogies Research Group (1995). *Fluid concepts &*

    creative analogies. New York: Basic Books.

Holyoak, K. J., & Thagard, P. (1995). *Mental leaps: Analogy in creative thought*.

    Cambridge, MA: MIT Press.

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A

    theory of analogical access and mapping. *Psychological Review, 104*, 427-466.

Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational

    inference and generalization. *Psychological Review, 110*, 220-264.

Johnson, H. M., & Seifert, C. M. (1992). The role of predictive features in retrieving

    analogical cases. *Journal of Memory & Language*, 31, 648-667.

Keil, F. C. & Batterman, N. (1984). A characteristic-to-defining shift in the acquisition

    of word meaning. *Journal of Verbal Learning and Verbal Behavior, 23*, 221-

    236.

Keller, J., & Loewenstein, J. (in press). The cultural category of cooperation: A Cultural Consensus Model analysis for China and the US. *Organization Science*.

Koestler, A. (1969). *The act of creation*. New York: Macmillan.

Kolodner, J. (1993). *Case-based reasoning*. San Francisco, CA: Morgan Kaufmann.

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

Kurtz, K. J., & Loewenstein, J. (2007). Converging on a new role for analogy in problem solving and retrieval: When two problems are better than one. *Memory & Cognition*, *35*(2), 334-341.

Lewandowsky, S., Kalish, M., & Griffiths, T. L. (2000). Competing strategies in categorization: Expediency and resistance to knowledge restructuring. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*(6), 1666-1684.

Loewenstein, J., & Gentner, D. (2001). Spatial mapping in preschoolers: Close comparisons facilitate far mappings. *Journal of Cognition and Development*, 2(2), 189-219.

Loewenstein, J., & Gentner, D. (2005). Relational language and the development of relational mapping. *Cognitive Psychology*, 50, 315-363.

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

Loewenstein, J., Thompson L., & Gentner, D. (2003). Analogical learning in

negotiation teams: Comparing cases promotes learning and transfer.

Academy of Management Learning and Education, 2(2), 119-127.

Markman, A.B. & Ross, B. H. (2003). Category use and category learning.

Psychological Bulletin, 129(4), 592-613.

Markman, A.B., Taylor, E., & Gentner, D. (2007). Auditory presentation leads to

better analogical retrieval than written presentation. *Psychonomic Bulletin

and Review, 14*(6), 1101-1106

Markman, A. B., Wood, K. L., Linsey, J. S., Murphy, J. T., & Laux, J. P. (2009).

Supporting innovation by promoting analogical reasoning. In A. B. Markman

& K. L. Wood (Eds.), *Tools for innovation: The science behind the practical

methods that drive new ideas* (pp. 85-103). New York: Oxford University

Press.

Mason, L. (2007). Introduction: Bridging the cognitive and sociocultural approaches

in research on conceptual change: Is it possible? *Educational Psychologist,

42*(1), 1-7.

Medin, D. L., & Ross, B. H. (1989). The specific character of abstract thought:

Categorization, problem-solving, and induction. In R. J. Sternberg (Ed.),

Advances in the psychology of human intelligence (Vol. 5, pp. 189-223).

Hillsdale, NJ: Erlbaum.

Mednick, S. A. (1962). The associative basis of the creative process. *Psychological

Review, 69*(3), 220-232.

Moss, J., Kotovsky, K., & Cagan, J. (2007). The Influence of Open Goals on the

> Acquisition of Problem-Relevant Information. *Journal of Experimental*
>
> *Psychology: Learning, Memory and Cognition*, 33, 876-891.

Myles-Worsley, M., Johnston, W. A., & Simons, M. A. (1988). The influence of

> expertise on x-ray image processing. *Journal of Experimental Psychology:*
>
> *Learning, Memory, and Cognition, 14*(3), 553-557.

Novick, L. (1988). Analogical transfer, problem similarity, and expertise. *Journal of*

> *Experimental Psychology: Learning, Memory and Cognition, 14*, 510-520.

Paletz, S. B. F., & Schunn, C. D. (in press). A social-cognitive framework of

> multidisciplinary team innovation. *Topics in Cognitive Science.*

Reeves, L. M., & Weisberg, R.W. (1994). The role of content and abstract information

> in analogical transfer. *Psychological Bulletin*, 115, 381-400.

Ross, B. H. (1984). Remindings and their effects in learning a cognitive skill.

> *Cognitive Psychology, 16,* 371-416.

Ross, B. H. (1987). This is like that: The use of earlier problems and the separation of

> similarity effects. *Journal of Experimental Psychology: Learning, Memory &*
>
> *Cognition, 13*(4), 629-639.

Ross, B. H. (1989). Distinguishing types of superficial similarities: Different effects

> on the access and use of earlier problems. *Journal of Experimental Psychology:*
>
> *Learning, Memory and Cognition, 15*(3), 456-468.

Ross, B. H., & 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., & Kilbane, M. C. (1997). Effects of principle explanation and superficial similarity on analogical mapping in problem solving. *Journal of Experimental Psychology: Learning, Memory, & Cognition 23*(2), 427-440.

Schank, R. (1982). *Dynamic Memory: A theory of reminding and learning in computers and people.* Cambridge: Cambridge University Press.

Schwartz, D. L. (1995). The emergence of abstract representations in dyad problem solving. *The Journal of the Learning Sciences, 4*(3), 321-354.

Seifert, C. M., McKoon, G., Abelson, R. P., & Ratcliff, R. (1986). Memory connections beween thematically similar episodes. *Journal of Experimental Psychology: Learning, Memory and Cognition, 12*(2), 220-231.

Seifert, C. M.. Meyer, D. E., Davidson, N., Patalano, A. L., & Yaniv, I. (1995) Demystification of cognitive insight: Opportunistic assimilation and the prepared-mind hypothesis. In R. J. Sternberg & J. E. Davidson (Eds.), *The Nature of Insight* (pp. 65-124). Cambridge, MA: MIT Press.

Simon, H. A. (1947). *Administrative behavior*. New York: Macmillan.

Smith, C. L. (2007). Bootstrapping processes in the development of students' commonsense matter theories: Using analogical mappings, thought experiments, and learning to measure to promote conceptual restructuring. *Cognition and Instruction, 25*(4), 337-398.

Smith, C. L., Solomon, G. E. A., & Carey, S. (2005). Never getting to zero: Elementary school students' understanding of the infinite divisibility of number and matter. *Cognitive Psychology, 51*, 101-140.

Stern, E. (2005). Knowledge restructuring as a powerful mechanism of cognitive

    development: How to lay an early foundation for conceptual understanding

    in formal domains. In P. D. Tomlinson, J. Dockrell & P. Winne (Eds.),

    Pedagogy—teaching for learning (British Journal of Educational Psychology

    Monograph Series II, No.3) (pp. 153-169). Leicester: British Psychological

    Society.

Thompson, L., Gentner, D., & Loewenstein, J. (2000). Avoiding missed opportunities

    in managerial life: Analogical training more powerful than individual case

    training. *Organizational Behavior and Human Decision Processes*, 82(1), 60-

    75.

Tomasello, M. (1999). *The cultural origins of human cognition*. Cambridge, MA:

    Harvard University Press.

Trope, Y., & Liberman, N. (2003). Temporal construal. *Psychological Review, 110*,

    403–421.

Tulving, E. (1983). *Elements of episodic memory*. New York: Oxford University Press.

Tulving, E., & Thompson, D. M. (1973). Encoding specificity and retrieval processes

    in episodic memory. *Psychological Review*, 80(5), 352-373.

Vosniadou, S., & Brewer, W. F. (1987). Theories of knowledge restructuring in

    development. *Review of Educational Research, 57*(1), 51-67.

Ward, T. B. (2009). ConceptNets for flexible access to knowledge. In A. B. Markman

    & K. L. Wood (Eds.), *Tools for innovation: The science behind the practical

    methods that drive new ideas* (pp. 153-170). New York: Oxford University

    Press.

Weisberg, R. W. (1993). *Creativity: Beyond the myth of genius.* New York: Freeman.

Wharton, C. M., Holyoak, K. J., Downing, P. E., Lange, T. E., Wickens, T. D., & Melz, E. R.

(1994). Below the surface: Analogical similarity and retrieval competition in

reminding. *Cognitive Psychology, 26*, 64-101.

Whitehead, A. N. (1929). *The aims of education*. New York: Macmillian.

Yamauchi, T., & Markman, A.B. (2000). Inference using categories. *Journal of*

Experimental Psychology: Learning, Memory, and Cognition, 26, 776-795.