Analogical Reasoning, Psychology of

Intermediate article

Dedre Gentner, Northwestern University, Evanston, Illinois, USA

CONTENTS

Analogical reasoning is a kind of reasoning that applies between specific exemplars or cases, in which what is known about one exemplar is used to infer new information about another exemplar. The basic intuition behind analogical reasoning is that when there are substantial parallels across different situations, there are likely to be further parallels.

INTRODUCTION

Analogical thinking is ubiquitous in human cognition. First, analogies are used in explaining new concepts. Domains such as electricity or molecular motion, which cannot directly be perceived, are often taught by analogy to familiar concrete domains such as water flow or billiard-ball collisions. Within cognitive science, mental processes are likened to computer programs (e.g. neural networks; parallel or serial processes), or to searching within a space (e.g. mental distance; close or far associates). Such analogies can then serve as mental models to support reasoning in new domains. Another use of analogy is in making predictions within domains. When the stock market plunged in 2001 after the attack on the World Trade Center, many newswriters made an analogy to the 1929 Wall Street crash, and argued on this basis that the market would be higher after a few years (or that, because key causal conditions are different, the reverse would occur). Analogy is also important in creativity and scientific discovery, as discussed later. Finally, analogy is used in communication and persuasion. For example, environmentalists have compared the earth to Easter Island, where overpopulation and exploitation of the island's once-rich ecology led first to massive loss of species, and then to famine and societal collapse. Such persuasive analogies are meant to invite new inferences: for example, that continued population growth will lead to irreversible ecological decline.

Analogical processing involves several subprocesses. First, given a current topic, *analogical retrieval* is the process of being reminded of a past situation from long-term memory. Once two cases are present in working memory (either because of an analogical retrieval or simply through encountering two cases together), *analogical mapping* can occur. We will begin by discussing the mapping processes.

MAPPING AND USE

History

Important early work on analogical mapping came out of philosophy, notably Hesse's analysis of analogical models in science. Early psychological research on analogy focused on simple four-term analogies of the form a:b::c:d. In the 1970s and 1980s, artificial intelligence researchers introduced a new level of representational complexity and computational specificity. Patrick Winston explored computational algorithms for analogical matching and inference, and Jaime Carbonell modeled the transfer of solution methods from one problem to another. This kind of work inspired psychologists to lay out detailed process models of how analogies are represented and processed. The ensuing period has seen intense computational and psychological research, theory revision, and an expansion of the phenomena studied. The field of analogy continues to be characterized by extremely fruitful interchange between computational models and psychological research. (*See* **Analogy-making, Computational Models of**)

Analogical Mapping

Analogical mapping is the core process in analogy. In a typical instance of analogical mapping, a familiar situation – the *base* or *source* description – is matched with a less familiar situation – the *target* description. The familiar situation suggests ways of viewing the newer situation as well as further inferences about it. Analogical mapping requires *aligning* the two situations – that is, finding the correspondences between the two representations – and *projecting inferences* from the base to the target. Then the reasoner must *evaluate* the analogical match and its inferences. Two further processes that can occur are *re-representation* of one or both analogs to improve the match, and *abstraction* of the structure common to both analogs.

Structure-mapping theory (Gentner, 1983) aims to capture the psychological processes that carry out analogical mapping. According to this theory, the comparison process involves finding an alignment between the base and target representations that reveals common relational structure. On the basis of this alignment, further inferences are projected from base to target. People prefer to find an alignment that is *structurally consistent*: that is, there should be a *one-to-one correspondence* between elements in the base and elements in the target, and the arguments of corresponding predicates must also correspond (*parallel connectivity*). For example, in the analogy below, Timmy in (a) could be put in correspondence with Timmy in (b) (on the basis of a local entity match) or with Fluffy in (b) (on the basis of matching relational roles). People appear to entertain both possibilities during processing, but to settle on one or the other by the end of the process.

(a) Lassie rescued Timmy.
(b) Timmy rescued Fluffy.

Another important early theory was Holyoak's (1985) *pragmatic mapping theory*, which focused on the use of analogy in problem-solving and held that analogical mapping processes are oriented towards attaining goals (such as solutions to problems). According to pragmatic mapping theory, it is goal relevance that determines what is selected in analogy. Holyoak and Thagard (1989) later combined this pragmatic focus with structural factors in their multi-constraint approach to analogy.

Analogical inference projection is a crucial part of the mapping process. Once an alignment is achieved, further inferences can be made by projecting information from the base (or source) domain into the target domain. For example, in the above analogy, suppose we knew more about event (a), such as:

(a) Lassie rescued Timmy because she loves him. She has beatiful brown eyes.
(b) Timmy rescued Fluffy.

In this case, the likely inference in (b) is that Timmy rescued Fluffy because he loves Fluffy. This ability to invite new inferences is central to analogy's role in reasoning. Importantly, analogical inference is rather selective. For example, we are unlikely to make the inference here that Timmy has brown eyes (or that he has four legs, even if we also know this to be true of Lassie).

This illustrates the *selection problem* in theories of analogical inference. If people projected everything known about the base into the target, analogy would be useless in reasoning. Fortunately, people do not do this. Thus a central aim of theories of analogy is to characterize this selection process. At least three factors have been proposed.

Holyoak and his colleagues have emphasized *goal relevance*: the inferences projected are those that fit with the reasoner's current goals in problem-solving.

A second factor, proposed in structure-mapping theory, is relational connectivity – more specifically, *systematicity*: a preference for projecting from matching systems of relations connected by higher-order relations such as *cause*, rather than projecting local matches. In many cases, goal relevance and systematicity will make the same predictions, because problem-solving goals often involve a focus on causal systems.

A third factor in selecting inferences, proposed by Keane, is *adaptability*: the ease with which a possible inference can be modified to fit the target.

There is evidence for all three of these criteria. Spellman and Holyoak (1996) showed that when two possible mappings are available for a given analogy, people will select the mapping whose inferences are relevant to their goals. Evidence for systematicity comes from the finding that when people read analogous passages and make inferences from one to the other, they are more likely to import a fact from the base to the target when it is causally connected to other matching predicates (Clement and Gentner, 1991; Markman, 1997). Finally, Keane (1996) found evidence that the degree of adaptability predicts which inferences are made from an analogy.

There remain many open questions. For example, according to the structure-mapping account, many different higher-order relations can provide inferential selection – including causal relations, deontic relations such as permission and obligation, and spatial relations such as symmetry and transitive increase. The challenge then is to delineate the set of higher-order relations that can serve this purpose. Another open question is the time course of these constraints. For example, do goals have special priority *during* the analogical mapping process, or do the effects of goals occur through influencing the initial representations of the two analogs (*before* the mapping process) or through selecting among multiple possible interpretations (*after* the mapping process)?

Evaluation

Once the common alignment and the candidate inferences have been discovered, the analogy is evaluated. *Evaluating* an analogy involves at least three kinds of judgment: (1) *structural soundness*: whether the alignment and the projected inferences are structurally consistent; (2) *factual correctness*: whether the projected inferences are false, true, or indeterminate in the target; and (3) *relevance*: whether the analogical inferences are relevant to the current goals. In practice, the relative importance of these factors varies quite a bit. In domains where little is known, or where there is disagreement about the facts – for example, in politics – goal relevance may be more important than factual correctness.

Abstraction

In *analogical abstraction*, the common system that represents the interpretation of an analogy is extracted and stored. This kind of schema abstraction helps to promote transfer to new exemplars. When people are asked to compare two analogous passages, they are better able to later retrieve and use their common structure (given a relationally similar probe) than are people who were given only one of the stories (Gick and Holyoak, 1983). Further studies have shown that actively comparing two analogous passages leads to better subsequent retrieval than reading the two passages separately. These findings are consistent with the claim that analogical alignment promotes the common structure and makes it more available for later use.

Analogy in Real-world Reasoning

Analogy is often used in common-sense reasoning to provide plausible inferences. It must be noted that analogy is not a deductive process. There is no guarantee that the inferences from a given analogy will be true in the target, even if the analogy is carried out perfectly and all of the relevant statements are true in the base. However, the set of implicit constraints described above make analogy a relatively 'tight' form of inductive reasoning. This may be why analogy is heavily used in arenas such as law, where clear reasoning is important but formal principles are often not sufficient to decide issues.

The lack of deductive certainty in analogical reasoning has a positive side. It means that analogy can suggest genuinely new hypotheses, whose truth could not be deduced from current knowledge. One arena in which this kind of analogical inferencing has been extensively studied is scientific reasoning and discovery. Nancy Nersessian has examined the role of analogy and other model-based reasoning processes in the thought processes of Faraday and Maxwell. Paul Thagard has discussed analogy as a contributor to conceptual revolutions in science. Kevin Dunbar has observed scientists in working microbiology laboratories and has found that analogy plays a large role in the discovery process.

Analogy appears to be very important in children's thinking, as Halford, Goswami, and others have argued. Children often use analogies from known domains as a way to fill in gaps in their knowledge of other domains. For example, Inagaki and Hatano (1987) asked five-year-old children hypothetical questions like 'What would happen if a rabbit were continually given more water?' The children often answered by using an analogy to humans: for example, 'I would get sick if I kept drinking water, and I think the rabbit would too.' Interestingly, children's answers tended to be more accurate when they used such analogies than when they did not. Children were most likely to use analogies to humans when the target was somewhat similar to humans, suggesting that for children (as for adults) similar analogs are more likely to be retrieved and are easier to align with the target than dissimilar analogs.

FACTORS THAT INFLUENCE ANALOGICAL MAPPING AND USE

People's fluency in carrying out analogical mappings is influenced by three broad kinds of factors. First are factors internal to the analogical mapping itself, such as *systematicity* – whether the common relational system possesses higher-order connective structure – and *transparency* – the degree to which corresponding elements are similar. The second category includes characteristics of the reasoner, such as age and expertise. The third includes task factors such as processing load, time pressure, and context.

Transparency and systematicity have been found to be important in analogical problem-solving. The transparency of the mapping between two analogous algebra problems – that is, the similarity between corresponding objects – is a good predictor of people's ability to notice and apply solutions from one problem to the other. For example, Ross (1989) taught people algebra problems and later gave them new problems that followed the same principles. People were better able to map the solution from a prior problem to a current problem when the corresponding objects were highly similar between the two problems: for example, 'How many golf balls per golfer' → 'How many tennis balls per tennis player'. They performed worst in the *cross-mapped* condition, in which similar objects appeared in different roles across the two problems: for example, 'How many golf balls per golfer' → 'How many tennis players per tennis ball'.

The intrinsic factors of transparency and systematicity interact with characteristics of the reasoner, notably age and experience. Gentner and Toupin (1986) gave children a simple story and asked them to re-enact the story with new characters. Both six- and nine-year-olds performed far better when the corresponding characters were highly similar between the two stories than when they were different, and they performed worst when similar characters played different roles across the two stories (the *cross-mapped* condition). Thus both age groups were sensitive to the transparency of the correspondences. In addition, older children (but not younger children) benefited from systematicity – that is, from hearing a summary statement that provided an overarching social or causal moral.

The developmental change found here is an instance of the *relational shift*: a shift from focusing on object matches to focusing on relational matches. Some researchers have suggested that this shift is driven by gains in knowledge (Gentner and Rattermann, 1991), while others propose that it results from a developmental increase in processing capacity (Halford, 1993).

The third class of factors affecting analogical processing concerns task variables such as time pressure, processing load, and immediate context. One generalization that emerges from several studies is that making relational matches requires more time and processing resources than making object–attribute matches. For example, Goldstone and Medin (1994) found that when people are forced to terminate processing early, they are strongly influenced by local attribute matches (such as *A* with *A* in the example below), even in cases where they would choose a relational match (such as *A* with *P*) if given sufficient time:

A above M and *P above A*

Adult performance in mapping tasks is also influenced by immediately preceding experiences. For example, in the *one-shot mapping task* (Markman and Gentner, 1993) subjects are shown a pair of cross-mapped pictures, such as *a robot repairing a car* and *a man repairing a robot*. The experimenter points to the robot in the first picture and the subject indicates which object in the second picture 'goes with' it. Subjects often choose the object match (e.g. the other robot). However, if they have previously rated the similarity of the pair, they are likely to choose the relational correspondence (the repairman). These findings suggest that carrying out a similarity comparison induces a structural alignment.

Kubose, Holyoak, and Hummel used this one-shot mapping task to show that processing load influences analogical processing. The experimenter pointed to the cross-mapped object in the first picture (the robot) and subjects were instructed to point to the relational correspondence (the repairman) in the second picture. Subjects made more object-mapping errors when given an extra processing load, such as having to count backwards. Recent work by Holyoak and his colleagues also suggests that damage to the prefrontal cortex is associated with detriments in analogical tasks, although it is not clear whether this results from specific involvement of the prefrontal cortex in analogical processing or from more general factors such as inhibitory control.

Summary

Research on analogical mapping has revealed a set of basic phenomena that characterize human analogical processing (see Table 1). A striking feature of analogical mapping is the importance of systematic, structurally connected representations. Commonalities that are interconnected into a relational system are considered to be more central to a comparison than are those that are not. Connected systems are easier to map to a new domain than are unconnected sets, leading to better transfer in analogy and problem-solving. Systematicity also governs inferences: inferences are projected from interconnected systems in the base to fill out corresponding structure in the target. Even the

Table 1. Basic phenomena of analogy (adapted from Gentner and Markman, 1995, 1997; see also Hummel and Holyoak, 1997)

1 *Relational similarity*	Analogies involve relational commonalities; object commonalities are optional.
2 *Structural consistency*	Analogical mapping involves one-to-one correspondence and parallel connectivity.
3 *Systematicity*	In interpreting analogy, connected systems of relations are preferred over isolated relations.
4 *Candidate inferences*	Analogical inferences are generated via structural completion.
5 *Alignable differences*	Differences that are connected to the common system are rendered more salient by a comparison.
6 *Interactive mapping*	Analogy interpretation depends on both terms. The same term yields different interpretations in different pairings.
7 *Multiple interpretations*	Analogy allows multiple interpretations of a single comparison.
8 *Cross-mapping*	Analogies are more difficult to process when there are competing object matches.

differences associated with a similarity comparison are influenced by systematicity: the differences that are psychologically salient in a comparison are those that are connected to the common system. In addition, goal relevance may have effects over and above the effects of connected relational structure.

RETRIEVAL OF ANALOGS

So far, we have discussed the processing of an analogy once both analogs are present. When we turn to the issue of what leads people to think of analogies, we see a very different pattern of results. People often fail to retrieve potentially useful analogs, even when there is an excellent structural match, and even when they clearly have retained the material in memory. For example, Gick and Holyoak (1983) gave subjects a thought problem: how to cure an inoperable tumor without using a strong beam of radiation that would kill the surrounding flesh. Only about 10 percent of the participants came up with the ideal solution, which is to converge on the tumor with several weak beams of radiation. If given a prior analogous story in which soldiers converged on a fort, three times as many people (about 30 percent) produced convergence solutions. Yet the majority of participants still failed to think of the convergence solution. Surprisingly, when these people were simply told to think back to the story they had heard, the percentage of convergence solutions again tripled, to 80–90 percent. Thus, the fortress story had been retained in memory, but it was not retrieved by the analogous tumor problem. The implication is clear. Even when a prior experience has been successfully stored in memory, it might not be retrieved when a person encounters a new analogous situation where it would be useful.

When we ask what does facilitate analogical retrieval, one major factor emerges: the similarity between the analogs. As noted earlier, similarity is one of the factors that facilitates analogical mapping; but it has a much larger effect on retrieval. For example, Gentner *et al.* (1993) gave subjects a set of stories to remember and later showed them probe stories that were either surface-similar to their memory item (e.g. similar objects and characters) or structurally similar (i.e. analogous, with similar higher-order causal structure). Surface similarity was the best predictor of whether people would be reminded of the prior stories; people were two to five times more likely to retrieve prior stories with only surface commonalities than with only structural commonalities. However, their judgments of the goodness of the match were completely different. They rated the surface-similar pairs (including their own remindings) as low in inferential value and in similarity, and preferred the structurally similar pairs. A similar dissociation between reminding and use has been found in problem-solving tasks: remindings of prior problems are strongly influenced by surface similarity, even though structural similarity better predicts success in solving current problems (Ross, 1989). This failure to access potentially useful analogs (unless they are highly similar to the target) is an instance of the *inert knowledge* problem in education. One piece of good news is that it appears that domain expertise may improve matters somewhat. For example, Novick (1988) found that people with mathematics training retrieved fewer surface-similar lures in a problem-solving task than did novice mathematicians. Moreover, experts were quicker to reject these false matches than were novices.

One factor that may contribute to experts' success in analogical retrieval is *representational uniformity*: the extent to which the relations in the memory trace are represented similarly to those in the probe. Clement *et al.* (1994) explored the effect of relational predicate similarity on analogical access and mapping between stories. They found that retrieval was more likely when the probe and target contained synonymous terms (*manifest* similarity) than when they contained loosely similar predicates such as 'munched' and 'consumed' (*latent* similarity). However, unlike retrieval, analogical mapping when both situations were present was relatively unaffected by the latent–manifest distinction. In analogical reminding, with only the current situation in working memory, success depends on the degree of match of the pre-existing representations; whereas during mapping, with both situations present in working memory, there is opportunity for re-representation.

ANALOGICAL LEARNING

Analogical comparison can lead to new learning in at least four ways: analogical abstraction, inference projection, difference detection, and re-representation. The first two we have already discussed. In *analogical abstraction*, the structure common to base and target is noticed and extracted. Sometimes the common system is stored as a separate representation; this is referred to as *schema abstraction* (Gick and Holyoak, 1983). In *inference projection*, a proposition from the base is mapped to the target. If it is retained as part of the target structure, then learning has occurred. In *difference detection*, carrying out a comparison process leads people to notice certain differences – namely, those connected to the common structure. In *re-representation*, two non-identical predicates are aligned and decomposed (or abstracted) to find their commonalities, resulting in a re-representation that contains a common predicate: for example, comparing *chase*(dog, cat) and *follow*(detective, suspect) might result in *pursue* (entity1, entity2). A further kind of knowledge change, hypothesized to take place in scientific discovery, is *restructuring*, in which the target undergoes a radical change in structure.

COMPUTATIONAL MODELS

The interplay between computational models and psychological studies has been extremely productive in analogical research. Current models include Boicho Kokinov's AMBR model, which integrates retrieval and mapping; Keane's IAM model, which utilizes an incremental mapping algorithm; Halford's tensor product model; and the systems of Doug Hofstadter and his colleagues Melanie Mitchell and Robert French, which integrate perceptual processing with analogical matching. (*See* **Analogy-making, Computational Models of**)

Analogical modeling has made great strides, but there are still open questions. At present no model fully captures human analogy processing. Two challenges for analogical models are (1) the *selection problem* discussed above – namely, how to avoid indiscriminate inferencing; and (2) the problem of *representational flexibility* – that is, how to achieve a matching process that does not require absolute identity matches. Falkenhainer *et al.*'s (1989) SME, which uses a local-to-global alignment and inference process over structured symbolic representations, meets the benchmarks in Table 1 and can capture selective matching and inference, as well as schema abstraction. But it is not yet sufficiently flexible in its match process. Another leading model, Hummel and Holyoak's (1997) LISA model, uses a combination of distributed representations of concepts and structured representations of the relational connections (necessary for achieving structural consistency in mapping). It uses a connectionist temporal-binding algorithm and makes its matches in a serial order partly guided by the experimenter. LISA's use of distributed representations allows for flexible matching, and unlike most models of analogy, it attempts to capture working memory limitations. However, it has yet to solve the selectivity problem in inferencing.

THE FUTURE

Of the many research questions that remain, four stand out as particularly interesting and timely. First is the role of analogy in cognitive development: how much of children's rapid learning can be attributed to the processes of comparing and drawing inferences between partially similar situations? Second is tracing the neuropsychology of analogical processes: what areas of the brain are implicated, and what is the course of processing? Third is the exploration of analogy in animal cognition. Comparative research so far indicates that humans excel in analogical ability, yet this ability exists in certain other species as well – for example, in chimpanzees and dolphins. Cross-species comparisons may help us decompose the cognitive components of analogical ability. A final important research frontier is the integration of analogy into larger models of cognition.

References

Clement CA and Gentner D (1991) Systematicity as a selection constraint in analogical mapping. *Cognitive Science* **15**: 89–132.

Clement CA, Mawby R and Giles DE (1994) The effects of manifest relational similarity on analog retrieval. *Journal of Memory and Language* **33**: 396–420.

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

Gentner D (1983) Structure-mapping: a theoretical framework for analogy. *Cognitive Science* **7**(2): 155–170.

Gentner D and Markman AB (1995) Analogy-based reasoning in connectionism. In: Arbib MA (ed.) *The Handbook of Brain Theory and Neural Networks*, pp. 91–93. Cambridge, MA: MIT Press.

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

Gentner D and Rattermann MJ (1991) Language and the career of similarity. In: Gelman SA and Byrnes JP (eds) *Perspectives on Thought and Language: Interrelations in Development*, pp. 225–277. London, UK: Cambridge University Press.

Gentner D, Rattermann MJ and Forbus KD (1993) The roles of similarity in transfer: separating retrievability from inferential soundness. *Cognitive Psychology* **25**: 524–575.

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

Gick ML and Holyoak KJ (1983) Schema induction and analogical transfer. *Cognitive Psychology* **15**(1): 1–38.

Goldstone RL and Medin DL (1994) Time course of comparison. *Journal of Experimental Psychology: Learning, Memory and Cognition* **20**(1): 29–50.

Halford GS (1993) *Children's Understanding: The Development of Mental Models*. Hillsdale, NJ: Lawrence Erlbaum.

Holyoak KJ (1985) The pragmatics of analogical transfer. In: Bower GH (ed.) *The Psychology of Learning and Motivation: Advances in Research and Theory*, vol. 19, pp. 59–87. New York, NY: Academic Press.

Holyoak KJ and Thagard PR (1989) Analogical mapping by constraint satisfaction. *Cognitive Science* **13**(3): 295–355.

Hummel JE and Holyoak KJ (1997) Distributed representations of structure: a theory of analogical access and mapping. *Psychological Review* **104**(3): 427–466.

Inagaki K and Hatano G (1987) Young children's spontaneous personification as analogy. *Child Development* **58**: 1013–1020.

Keane MT (1996) On adaptation in analogy: tests of pragmatic importance and adaptability in analogical problem solving. *Quarterly Journal of Experimental Psychology* **49**: 1062–1085.

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

Markman AB and Gentner D (1993) Structural alignment during similarity comparisons. *Cognitive Psychology* **25**: 431–467.

Novick LR (1988) Analogical transfer, problem similarity, and expertise. *Journal of Experimental Psychology: Learning, Memory and Cognition* **14**: 510–520.

Ross BH (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**: 456–468.

Spellman BA and Holyoak KJ (1996) Pragmatics in analogical mapping. *Cognitive Psychology* **31**: 307–346.

Further Reading

Carbonell JG (1983) Learning by analogy: formulating and generalizing plans from past experience. In: Michalski RS, Carbonell JG and Mitchell TM (eds) *Machine Learning: An Artificial Intelligence Approach*, vol. 1, pp. 137–161. Palo Alto, CA: Tioga.

Dunbar K (1995) How scientists really reason: scientific reasoning in real-world laboratories. In: Sternberg RJ and Davidson JE (eds) *The Nature of Insight*, pp. 365–395. Cambridge, MA: MIT Press.

French R (1995) *The Subtlety of Sameness: A Theory and Computer Model of Analogy-making*. Cambridge, MA: MIT Press.

Gentner D, Holyoak KJ and Kokinov BN (eds) (2001) *The Analogical Mind: Perspectives from Cognitive Science*. Cambridge, MA: MIT Press.

Goswami U (1992) *Analogical Reasoning in Children*. Hillsdale, NJ: Lawrence Erlbaum.

Halford GS, Wilson WH and Phillips S (1998) Processing capacity defined by relational complexity: implications for comparative, developmental, and cognitive psychology. *Behavioral and Brain Sciences* **21**: 803–864.

Hesse MB (1966) *Models and Analogies in Science*. Notre Dame, IN: University of Notre Dame Press.

Hofstadter DR and Mitchell M (1994) The Copycat project: a model of mental fluidity and analogy-making. In: Holyoak KJ and Barnden JA (eds) *Advances in Connectionist and Neural Computation Theory*, vol. 2: *Analogical Connections*, pp. 31–112. Norwood, NJ: Ablex.

Holyoak KJ and Thagard PR (1995) *Mental Leaps: Analogy in Creative Thought*. Cambridge, MA: MIT Press.

Keane MT (1990) Incremental analogising: theory and model. In: Gilhooly KJ, Keane MTG, Logie RH and Erdos G (eds) *Lines of Thinking*, vol. 1. Chichester, UK: John Wiley.

Kokinov BN and Petrov AA (2001) Integrating memory and reasoning in analogy-making: the AMBR model. In: Gentner D, Holyoak KJ and Kokinov BN (eds) *The Analogical Mind: Perspectives from Cognitive Science*, pp. 161–196. Cambridge, MA: MIT Press.

Nersessian NJ (1984) *Faraday to Einstein: Constructing Meaning in Scientific Theories*. Dordrecht, Netherlands: Nijhoff.

Thagard P (1992) *Conceptual Revolutions*. Princeton, NJ: Princeton University Press.

Winston PH (1982) Learning new principles from precedents and exercises. *Artificial Intelligence* **19**: 321–350.