

Analogy and Qualitative Representations in the Companion Cognitive Architecture

Kenneth D. Forbus, Thomas Hinrichs

■ *The Companion cognitive architecture is aimed at reaching human-level AI by creating software social organisms — systems that interact with people using natural modalities, working and learning over extended periods of time as collaborators rather than tools. Our two central hypotheses about how to achieve this are (1) analogical reasoning and learning are central to cognition, and (2) qualitative representations provide a level of description that facilitates reasoning, learning, and communication. This article discusses the evidence we have gathered supporting these hypotheses from our experiments with the Companion architecture. Although we are far from our ultimate goals, these experiments provide strong evidence for the utility of analogy and qualitative representation across a range of tasks. We also discuss three lessons learned and highlight three important open problems for cognitive systems research more broadly.*

Every cognitive architecture starts with a set of theoretical commitments. We have argued (Forbus 2016) that human-level artificial intelligences will be built by creating sufficiently smart software social organisms. By that we mean systems capable of interacting with people using natural modalities, operating and learning over extended periods of time, as apprentices and collaborators, instead of as tools. Just as we cannot directly access the internal representations of the people and animals we work with, cognitive systems should be able to work with us on our terms. But how does one create such systems? We have two core hypotheses, inspired by research in cognitive science:

Our first core hypothesis is that *analogical reasoning and learning are central to human cognition*. There is evidence that processes described by Gentner's (1983) structure-mapping theory of analogy and similarity operate throughout human cognition, including visual perception (Sagi, Gentner, and Lovett 2012), reasoning and decision making (Markman and Medin 2002), and conceptual change (Gentner et al. 1997).

Our second core hypothesis is that *qualitative representations (QRs) are a key building block of human conceptual structure*. Continuous phenomena and systems permeate our environment and our ways of thinking about it. This includes the physical world, where qualitative representations have a long track record of providing human-level reasoning and performance (Forbus 2014), but also in social reasoning (for example, degrees of blame [Tomai and Forbus 2007]). Qualitative representations carve up continuous phenomena into symbolic descriptions that serve as a bridge between perception and cognition, facilitate everyday reasoning and communication, and help ground expert reasoning.

The focus of the Companion cognitive architecture (Forbus, Klenk, and Hinrichs 2009) is on higher-order cognition: conceptual reasoning and learning, and learning through interactions with others, by contrast with architectures that have focused on skill learning (for example, ACT-R [Anderson 2009] and SOAR [Laird 2012]). In Newell's (1990) time-scale decomposition of cognitive phenomena, conceptual reasoning and learning occur in what are called the *rational and social bands*,¹ unlike many architectures, which start with Newell's cognitive band (Laird, Lebiere, and Rosenbloom 2017). Thus we approximate subsystems whose operations occur at faster time scales, using methods whose outputs have reasonable cognitive fidelity, although we are not making theoretical claims about them. For example, in Companions constraint checking and simple inferences are carried out through a logic-based truth-maintenance system (Forbus and de Kleer 1994). Similarly, the Companion architecture is implemented as a multiagent system, capable of running on a single laptop or distributed across many cluster nodes, depending on the task. We suspect that many data-parallel operations and a number of coarse-grained parallel operations are important for creating robust software organisms, but this current organization is motivated more by engineering concerns. By contrast, we spend considerable effort on sketch understanding and natural language understanding, since interaction through natural modalities is a core concern.

Cognitive systems are knowledge rich (for example, McShane [2017]). The Companion architecture uses the Cyc² ontology and knowledge base (KB) contents, plus our own extensions including additional linguistic resources, representations for visual/spatial reasoning, and the hierarchical task network (HTN) plans that drive Companion operations. This is in sharp distinction with current practice in machine learning, where the goal is that systems for every task must be learned from scratch. The machine-learning (ML) approach can be fine for specific applications, but it means that many independently learnable factors in a complex task must be learned at the same time. This intermingling leads to requiring far more training data than people need, and reduces transfer of learned knowledge from one task to another. The cognitive-systems approach strives for cumulative learning, where an architecture that is general purpose can learn to do one task and successfully applies that knowledge to learn other tasks.

The remainder of this article summarizes the evidence provided by Companion experiments for our core hypotheses. We start by summarizing the ideas underlying our hypotheses, briefly introducing structure mapping and our models of its processes for analogy followed by the key aspects of qualitative reasoning. Then we describe the evidence that analogical and qualitative representations play a variety of

roles in cognition. We close with some lessons learned and open problems.

Analogical Reasoning and Learning

Dedre Gentner's (1983) structure-mapping theory proposed that analogy involves the construction of mappings between two structured, relational representations. These mappings contain *correspondences* (that is, what goes with what), *candidate inferences* (that is, what can be projected from one description to the other, based on the correspondences), and a score indicating the overall quality of the match. Typically, the *base* is the description about which more is known, and the *target* is the description about which one is trying to reason about, and hence inferences are made from base to target by default. Inferences in the other direction, *reverse candidate inferences*, can also be drawn, which is important for difference detection, discussed more later.

There is ample evidence for the psychological plausibility of structure mapping, out of Gentner's lab and others (Forbus 2001). This evidence includes some implications that may surprise AI researchers. Here are three such findings:

(1) The same computations that people use for analogy are also used for everyday similarity (Markman and Gentner 1993). Most AI researchers model similarity as the dot product of feature vectors or other distributed representations. But those approaches are not compatible with psychological evidence that indicates, even for visual stimuli, that relational representations are important (Goldstone, Medin, and Gentner 1991).

(2) Differences are closely tied to similarity (Markman and Gentner 1996). There is an interesting dissociation between difference detection and naming a difference: it is faster to detect that things are different when they are very different, but faster to name a difference when they are very similar (Sagi, Gentner, and Lovett 2012). This falls directly out of our computational model of analogical matching below.

(3) Analogical comparison, especially within-domain comparison, can happen unconsciously. For example, people reading a story can be pushed between two different interpretations of it, based on a different story that they read previously (Day and Gentner 2007). Moreover, they report being sure that the second story was complete in itself, and nothing from the previous story influenced them. This implies, for example, that in any of the System 1 / System 2 models (for example, Kahneman [2011]), analogy and similarity are actually used in both systems.

Our computational models of analogical processes have been designed both as cognitive models and as performance systems. That is, each of them has been used to explain (and sometimes predict) psychological findings, while also being used in AI systems

whose sole goal was achieving some new functionality. These models are SME, MAC/FAC, and SAGE.

The structure-mapping engine (SME) (Forbus et al. 2016) models analogical mapping. It computes up to three mappings, using a greedy algorithm to operate in polynomial time. It can also operate incrementally, extending mappings as new facts are added to the base or target, which can be useful in problem solving and language understanding.

Many are called but few are chosen (MAC/FAC) (Forbus, Gentner, and Law 1995). MAC/FAC models analogical retrieval, using two map/reduce stages. The MAC stage uses a nonstructured vector representation, automatically derived from structured representations, such that their dot products provide an estimate of how SME will score the original structural representations. The best, and up to two more if sufficiently close, are passed to the FAC stage, which uses SME in parallel to match the original structured representations. The MAC stage provides scalability, while the FAC stage provides the precision and inferences needed to support reasoning.

The sequential analogical generalization engine (SAGE) (McLure, Friedman, and Forbus 2015) models analogical generalization. For each concept, it maintains a generalization pool that is incrementally updated as new examples arrive. The pool contains generalizations, which are structured descriptions with a probability attached to each statement, and unassimilated examples. For each new example, the most similar item from the pool is retrieved through MAC/FAC. If sufficiently similar, they are assimilated. The assimilation process involves updating the probabilities for each statement, based on the contents of the match and the differences, and replacing any nonidentical entities with skolems. Thus commonalities become highlighted and accidental information becomes deprecated. Note that logical variables are not introduced, since generalizations can still be applied to new situations through analogy.

In computational terms, we think of these as an *analogy stack*: An organization of analogical processing for cognitive systems that we believe has wide applicability. In the Companion cognitive architecture, analogical operations are more primitive than back chaining. Rule-based reasoning and logical deduction are used for encoding, case construction, constraint checking, and simple inferences. The architecture also includes an AND/OR graph problem solver and an HTN planner, but all of these freely use the analogical processes above in their operations.

Qualitative Representations

People have robust mental models that help them operate in the world; for example, we cook, navigate, and bond with others. A common concern in these models is how to deal with continuous properties and phenomena. In reasoning about the physical

world, we must deal with quantities like heat, temperature, and mass, and arrange for processes like flows and phase changes to cook. Spatial factors are key in many physical situations, such as finding safe ground during a flood. In social reasoning, we think about quantities like degree of blame, how strong a friendship is, and the relative morality of choices we must make. Even in metacognition, we estimate how difficult problems are, and decide whether to carry on or give up on a tough problem. Qualitative representations provide more abstract descriptions of continuous parameters and phenomena, which are more in line with data actually available to organisms. We may not be able to estimate accurately how fast something is moving, but we can easily distinguish between rising and falling, for example. People routinely come to reasonable conclusions with very little information, and often do so rapidly. Combining qualitative representations with analogy provides an explanation for this (Forbus and Gentner 1997).

Broadly, qualitative representations can be decomposed into those concerned with quantities and those concerned with space. Quantities are an important special case because many continuous parameters are one dimensional, and this imposes additional constraints on reasoning. We use qualitative process theory (QP theory, Forbus [1984]) as our account of quantities and continuous processes. Spatial qualitative representations are typically grounded in metric representations (that is, the metric diagram/place vocabulary model [Forbus, Nielsen, and Faltings 1991]), which provide the functional equivalent of visual processing to support spatial reasoning. Once places are constructed, often purely qualitative reasoning can be done on them, through spatial calculi (Cohn and Renz 2008). Our current model of visual and spatial reasoning, CogSketch (Forbus et al. 2011), automatically produces cognitively plausible visual and spatial representations from vector inputs (for example, digital ink, copy/paste from PowerPoint), and is integrated into the Companion architecture.

Some Roles of Analogy and Qualitative Representations in Cognition

We provide evidence for our hypotheses by examining a range of tasks, showing how analogy and qualitative representations have been used to enable the Companion architecture to perform them.³

Reasoning

Textbook problem solving involves solving the kinds of problems that are posed to students learning science and engineering. In experiments conducted with the aid of the Educational Testing Service (ETS), we showed that a Companion, using representations generated by ETS and Cycorp, could learn to transfer knowledge through analogy to new problems across

six different within-domain conditions (Klenk and Forbus 2009). It operated by using MAC/FAC to retrieve worked solutions, which were at the level of detail found in textbooks, rather than the internals of the reasoning system, which neither ETS nor Cycorp had access to. Equations to solve new problems were projected as candidate inferences onto the new problem, with their relevance ascertained by overlapping event structure between the new and old problem. This included solving larger, more complex problems by combining analogies involving two solutions from simpler problems. Qualitative representations played three roles in this reasoning: First, some problems themselves are qualitative in their nature, such as asking whether something will move at all in a given situation, and if so, will its velocity and acceleration be constant or increasing. The second role was expressing the conditions under which an equation is applicable. The third role was translating real-world conditions into parameter values (for example, at the top of its trajectory, the velocity of a rising object is zero).

The role of qualitative representations for reasoning found in earlier stages of science education is even larger. Before high school, children are not asked to formulate equations, but to reason qualitatively, draw conclusions from graphs and tables, and tie knowledge learned in class to their everyday experience. An analysis of fourth-grade science tests suggests that qualitative representations and reasoning are sufficient for at least 29 percent of their questions (Crouse and Forbus 2016), and even more when qualitative spatial and visual representations are taken into account.

Decision Making

An example of decision making is MoralDM (Deghani et al. 2008), which models the influence of utilitarian concerns and sacred values on human decision making. An example of a sacred value is not taking an action that will directly result in someone's death, even if it would save more lives (see Scheutz [2017] for examples). MoralDM modeled the impact of sacred values through a qualitative order of magnitude representation, so that when a sacred value was involved, the utilitarian differences seemed negligible compared to violating that value. While it had some rules for detecting that a situation should invoke a sacred value, it also used analogy to apply such judgements from similar situations. Experiments indicate that increasing the size of the case library improved its decision, as measured against human judgments in independent experiments (Deghani et al. 2008), and using analogical generalization over cases improved performance further (Blass and Forbus 2015).

Commonsense Reasoning

Analogical reasoning may be at the heart of commonsense reasoning: Similar prior experiences can be used to make predictions but also provide explana-

tions for observed behavior (Forbus 2015). Moreover, even partial explanations relating aspects of a situation can be projected through candidate inferences, without a complete and correct logically quantified domain theory. For example, questions from a test of plausible reasoning can be answered by chaining multiple small analogies, similar to how rules are chained in deductive reasoning (Blass and Forbus 2016).

Understanding Language

Our research indicates that analogy can play at least four roles in natural language understanding. First, analogies are often explicitly stated in language, particularly in instructional materials. These need to be detected so that subsequent processing will gather up the base and target and the appropriate conclusions can be drawn. Analogical dialogue acts (Barbella and Forbus 2011) extend communication act theories with additional operations inspired by structure mapping, enabling a Companion to learn from such analogies expressed in natural language, as measured by improved query-answering performance. Second, for integrating knowledge learned by reading, it can be useful to carve up material in a complex text into cases. These cases can be used for answering comparison questions directly (Barbella and Forbus 2015), and for subsequent processing (for example, rumination [Forbus et al. 2007]). Third, analogy can also be used to learn evidence for semantic disambiguation, by accumulating cases about how particular word/word sense pairings were used in the linguistic context (including lexical, syntactic, and semantic information), and honed by SAGE to suggest how to disambiguate future texts (Barbella and Forbus 2013). Finally, analogy can be used to learn new linguistic constructions (McFate and Forbus 2016a), for example, understanding sentences like "Sue crutched Joe the apple so he wouldn't starve."

Our evidence to date suggests that QP theory provides a formalism for natural language semantics concerning continuous phenomena. Qualitative representations capture the level of detail found in linguistic descriptions, and mappings can be established from FrameNet representations (McFate and Forbus 2016b) to qualitative process theory. Moreover, qualitative representations can be learned by reading, with the knowledge produced improving strategy game performance (McFate, Forbus, and Hinrichs 2014).

Domain Learning

Learning from instruction often involves multiple modalities, for example, text and diagrams. Analogy can be used to fuse cross-modal information, by using SME to align the representations constructed for each modality and projecting information across them. For example, Lockwood (Lockwood and Forbus 2009) showed that semiautomatic language processing of a simplified English chapter of a book on basic

machines, combined with sketches for the diagrams, enabled a system to learn enough to correctly answer 12 out of 15 questions from the back of the chapter. Chang (2016) has shown that these ideas can work when the language processing is fully automatic. By using a corpus of instructional analogies aimed at middle-school instruction, expressed in simplified English and sketches, a Companion learned enough to correctly answer 11 out of 14 questions from the New York Regent's examination on the topics covered.

Companions have also been used to do cross-domain analogical learning, both in physical domains (Klenk and Forbus 2013) and in games (Hinrichs and Forbus 2011). For cross-domain learning, building up persistent mappings across multiple analogies, essentially making a translation table between domains, was important. In cross-domain game learning, SME was used to build up a metamapping between predicates in the domains, which then enabled the rest of the domain knowledge to be imported.

Qualitative representations provide a useful level of causal knowledge for learning the dynamics of complex domains (Hinrichs and Forbus 2012a), and can be used as a high-level representation for goals and strategies (Hinrichs and Forbus 2014). Qualitative reasoning supports reflective reasoning by treating goal activation levels as continuous quantities in order to represent mental states (Hinrichs and Forbus 2016). For example, in reasoning about a strategy game such as *Freeciv*,⁴ the learned qualitative model of domain dynamics enables identifying trade-offs, and high-level strategies such as first expanding one's empire, then building up its economy, can be expressed by processes that are carried forward by a player's actions.

Conceptual Change

One of the well-documented phenomena in cognitive science is conceptual change, that is, how people's mental models develop over long spans of time (for example, diSessa, Gillespie, and Esterly [2004]; Ioannides and Vosniadou [2002]). Catalogs of misconceptions have been created for a variety of phenomena, and in some cases, trajectories that learners take in mastering these concepts have been analyzed and documented. Friedman's assembled coherence theory of conceptual change (Friedman 2012) takes as its starting point qualitative, compositional model fragments, but initially tied to specific experiences. Analogical retrieval is used to determine which model fragments are retrieved and combined to create a model for a new situation. This theory explains both why mental models often seem very specific to particular classes of situations and yet, within those classes, are relentlessly consistent. Model fragments are constructed through SAGE and a set of transformation heuristics, and applied through MAC/FAC. Friedman's Companion-based TIMBER system has

successfully modeled the trajectories of intuitive force models (Friedman and Forbus 2010) using sketched behaviors, transitions between misconceptions when learning about the human circulatory system (Friedman and Forbus 2011), misconceptions about why there are seasons (Friedman, Forbus, and Sherin 2011), and debugging misconceptions about the day/night cycle by analogy (Friedman, Barbella, and Forbus 2012).

Cognitive State Sensing

Our mental life is governed in part by the ability to sense our internal state: in doing a forced-choice task, like the one in figure 1, often one choice "looks right" and we take it, even if we cannot always articulate why this is so. If it isn't obvious, we have to think about it more. We believe that the structural evaluation score computed by SME is a signal used in these computations. In simulating this task (Kandaswamy, Forbus, and Gentner 2014), if there is a strong difference between SME's scores for the two comparison, the highest score is chosen. But if not, rerepresentation techniques are used to modify the encodings of the stimuli (which are automatically generated through CogSketch) to attempt to discriminate between them. This model has captured the phenomena in several psychological experiments.

Emotions are another kind of cognitive state sensing, in our view. One property of emotions is that they are often rapid, but not always specific. Wilson, Forbus, and McLure (2013) suspect this is because appraisal information used in generating emotions is stored with memories and retrieved through analogy to produce a rapid response. This response is subsequently modified by slower cognitive processes, and eventually when new memories are consolidated, the appraisals are adjusted based on the outcome of the situation. Hence a problem that might have initially seemed scary, if mastered, gets stored as something doable. With this multiphase model of emotions, a Companion learning to solve problems was able to perform more effectively, and with emotional dynamics consistent with the literature.

Classification

The analogy stack described earlier can be used for classification, by doing analogical retrieval over the union of generalization pools representing the possible concepts of interest. The label associated with the pool where the best reminding came from serves as the identification of the concept, with the score providing an estimate of the quality of the classification and the correspondences of the mapping providing an explanation as to why that classification is appropriate. The candidate inferences provide further surmises about the new example. This method has been used to learn geospatial concepts in a strategy game (for example, isthmus, bay, and others), achieving 62 percent accuracy. By extending SAGE with automat-

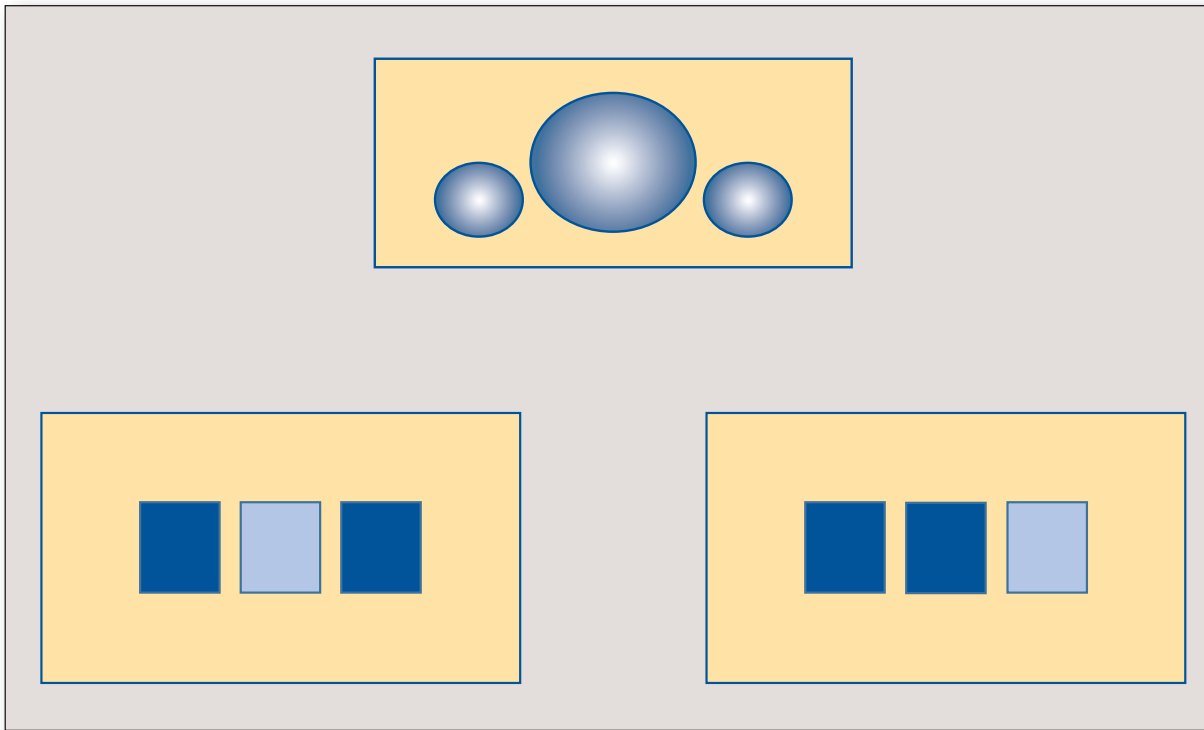


Figure 1. An Example of a Forced Choice Task.

ic near-miss detection, accuracy rose to 77 percent (McLure, Friedman, and Forbus 2015).

A Meta-Analysis

Analogy is more conservative than explanation-based learning, which generalizes after only one example, but requires many fewer examples than feature-based statistical learning systems — sometimes orders of magnitude fewer (for example, Kuehne, Gentner, and Forbus [2000]). Can we do even better by somehow combining ideas of structure-mapping with ideas from machine learning? Here is a meta-analysis that suggests a family of approaches: structure mapping is to relational representations what dot product is to feature vectors. This suggests that any machine-learning technique defined on feature vectors originally might be adapted to a hybrid method by replacing dot product with SME. For example, by combining a structured form of logistic regression with SAGE on the link plausibility/knowledge base completion task, Liang and Forbus (2015) showed that this approach can yield state-of-the-art results, with several orders of magnitude fewer examples, while also providing material for explanations. Another hybrid has been shown to produce state of the art performance on the Microsoft Paraphrase task (Liang et al. 2016). The advantage of using task-specific training is that similarity weights can be tuned to specific tasks, and learned negative weights (which SME does not produce) can improve discrimination. The drawback is that the representations learned, unless the weights are stored separately,

become task specific instead of general purpose. We believe this is a very useful approach that can be applied to many techniques, for example, using SME as a kernel for SVMs. But we know of no evidence that this is psychologically plausible.

Some Lessons Learned

We have learned a lot through these and other experiments. We summarize the three most important lessons here.

First, the models of structure-mapping processes turn out to be surprisingly robust. We have only made two extensions in tackling the span of tasks we have discussed. The first are *interim generalizations*, that is, a small set of generalization pools that exist in working memory, in contrast with normal generalization pools, which are held in long-term memory. Interim generalizations were motivated by both the rapid learning of children in experiments and by the unconscious story intrusion effects in adults. In the experiments listed, they are used in learning forced-choice tasks as a simple form of rerepresentation, that is, when an interim generalization is retrieved for an item, nonoverlapping aspects of the item are filtered out. The second are near misses. Winston's (1970) concept of near miss required a teacher to know the machine's single concept, and formulate an example that differed in only one way. McLure's observation (McLure, Friedman, and Forbus 2015) was that retrieving an item from a different generalization

pool over items in the intended label is, when the two labels are mutually exclusive, a near miss. He extended SAGE to construct positive and negative probabilistic hypotheses based on near misses, which are added to the retrieval criteria for the generalization pool to improve discriminability.

Second, type-level qualitative representations (Hinrichs and Forbus 2012b) are important for scalability and for flexible reasoning. In traditional qualitative reasoning, a system might start with a schematic or scenario description, and then formulate and use models to reason about it. By contrast, consider playing a strategy game like Freeciv, where constructing a civilization is the goal. A successful civilization involves a dozen or more cities, each with many parameters, and even more units. Formulating and maintaining explicit propositional qualitative models for each city, much less building a complete snapshot of the game state, can be prohibitively expensive. Moreover, it is often necessary to reason about planned units and cities that don't exist yet. Type-level representations make these problems simpler, although still not easy. They are amenable to analogy (because they are quantifier free) and also useful in natural language semantics: The QP frame representations we generate are neutral with respect to whether they are specific or generic, leaving this decision to later semantic processing, which has the contextual information needed to determine this.

Finally, we have found the Cyc ontology and KB contents to be extremely valuable. We developed our own reasoning engine because we needed source code access and, for us, analogy operations are central. The breadth and depth of the Cyc ontology has saved us countless hours, by letting us use existing representations instead of building our own. Even when our requirements diverge from Cyc's ontology, small additions usually suffice to bridge the gap. For example, we treat quantities as fluents, as opposed to specific values, so a single logical function suffices to translate⁵. Moreover, we have found microtheories to be an excellent way to impose modularity on KB organization and support more flexible reasoning.

Every operation is done with respect to a logical environment, for example, a microtheory and those it inherits from, thereby controlling what knowledge is used in a computation. Alternatives and qualitative states are represented as microtheories, as are subsketches within a sketch. In our reasoning engine, microtheories are also cases, which makes analogical reasoning even more tightly integrated.

Discussion and Open Problems

Our experience with Companions to date has provided strong evidence for our core hypotheses. Analogical processing and qualitative representations are useful for textbook problem solving, moral decision making, and commonsense reasoning. Analogy can play multiple roles in natural language understanding and visual problem solving, and qualitative representations provide a natural level of expressiveness for semantics and reasoning. Together they can successfully model a variety of phenomena in conceptual change, and provide a means of fusing information across modalities. They can be used inside cognitive systems for cognitive state sensing, to improve problem solving. They can be used for classification tasks. And there is now evidence that, by combining analogical processing with ideas from standard machine learning, state of the art performance on tasks of interest to the machine learning community can be attained, sometimes with orders of magnitude fewer examples and increased explainability.

To be sure, many open problems remain. We outline three here. The first is scaling. How much knowledge might be needed to achieve human-level AI? Aside from the initial endowment of knowledge, how much experience might be needed? A back of the envelope calculation (Forbus, Liang, and Rabkina 2017) suggests that a lower bound of at least 2 million generalization pools, built up by at least 45 million examples, would be needed to approximate some reasonable fraction of what a college graduate knows. For Companions, we will need to exploit the data-parallel capabilities built into

our analogical models through hardware, to reach this scale and maintain real-time performance. The second is that being able to use natural modalities to interact with people, in a fluent and adaptive manner, is a key bottleneck for progress. While for Companions simplified English syntax has been very productive, we would prefer them to be able to puzzle through more complex syntax. Our current hope is that learning linguistic constructions through analogy will enable this, but that is far from clear. The final issue is longevity, that is, creating cognitive systems that successfully perform while learning over weeks, months, and years at a time. Most cognitive architectures, including ours, are fired up for experiments, the experiments are run, and then they are shut down. Achieving robust learning while maintaining accurate performance, even with people inspecting the system's internals, is quite challenging (Mitchell et al. 2015). Enabling Companions to build up solid internal models of their own reasoning and learning processes seems to be the most promising approach to us. We think that the notions of attention and control that Bello and Bridewell (2017) explore will be an important part of the solution to longevity.

Acknowledgments

We thank the Office of Naval Research, the Air Force Office of Scientific Research, the National Science Foundation, and the Defense Advanced Research Projects Agency for their support through multiple programs. Statements and opinions expressed may not reflect the position or policy of the United States government, and no official endorsement should be inferred. David Barbella, Maria Chang, Morteza Dehghani, Scott Friedman, Matthew Klenk, Kate Lockwood, Andrew Lovett, Joe Blass, Max Crouse, Subu Kandaswamy, Chen Liang, Clifton McFate, Matthew McLure, Irina Rabkina, and Madeline Usher contributed heavily to this research.

Notes

1. In models of some processes, for example, analogical matching, we capture phenomena at the Cognitive band as well.
2. www.cyc.com

3. Here we stick to models that use the entire Companion cognitive architecture. Other models indicate that these same ideas are relevant in modeling visual analogies, including geometric analogies (Lovett et al. 2009b), oddity tasks (Lovett and Forbus 2011), and Ravens' Progressive Matrices (Lovett et al. 2017).

4. www.freeciv.org.

5. The logical function QPQuantityFn we introduced has as its domain Cyc continuous quantities and has as its range QP theory fluents.

References

- Anderson, J. R. 2009. *How Can the Human Mind Occur in the Physical Universe?* Oxford, UK: Oxford University Press.
- Barbella, D., and Forbus, K. 2011. Analogical Dialogue Acts: Supporting Learning by Reading Analogies in Instructional Texts. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence* (AAAI 2011). Palo Alto, CA: AAAI Press.
- Barbella, D., and Forbus, K. 2013. Analogical Word Sense Disambiguation. *Advances in Cognitive Systems* 2: 297–315.
- Barbella, D., and Forbus, K. 2015. Exploiting Connectivity for Case Construction in Learning by Reading. *Advances in Cognitive Systems* 4: 169–186.
- Bello, P., and Bridewell, W. 2017. There Is No Agency Without Attention. *AI Magazine* 38(4). doi.org/10.1609/aimag.v38i4.2742
- Blass, J., and Forbus, K. D. 2015. Moral Decision-Making by Analogy: Generalizations Versus Exemplars. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*. Palo Alto, CA: AAAI Press.
- Blass, J. A., and Forbus, K. D. 2016. Modeling Commonsense Reasoning via Analogical Chaining: A Preliminary Report. In *Proceedings of the 38th Annual Meeting of the Cognitive Science Society*, Philadelphia, PA, August. Austin, TX: Cognitive Science Society Inc.
- Chang, M. D. 2016. Capturing Qualitative Science Knowledge with Multimodal Instructional Analogies. Doctoral dissertation, Northwestern University, Department of Electrical Engineering and Computer Science, Evanston, Illinois.
- Cohn, A., and Renz, J. 2008. Qualitative Spatial Representation and Reasoning. In *Handbook of Knowledge Representation*, ed. F. van Harmelen, V. Lifschitz, B. Porter, 551–596. Amsterdam, The Netherlands: Elsevier. doi.org/10.1016/S1574-6526(07)03013-1
- Crouse, M., and Forbus, K. 2016. Elementary School Science as a Cognitive System Domain: How Much Qualitative Reasoning Is Required? *Advances in Cognitive Systems* 4.
- Day, S. and Gentner, D. 2007. Nonintentional Analogical Inference in Text Comprehension. *Memory and Cognition* 35(1): 39–49. doi.org/10.3758/BF03195940
- Dehghani, M.; Tomai, E.; Forbus, K.; and Klenk, M. 2008. An Integrated Reasoning Approach to Moral Decision-Making. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence* (AAAI). Palo Alto, CA: AAAI Press.
- diSessa, A.; Gillespie, N.; and Esterly, J. 2004. Coherence versus Fragmentation in the Development of the Concept of Force. *Cognitive Science* 28: 843–900. doi.org/10.1207/s15516709cog2806_1
- Forbus, K. 1984. Qualitative Process Theory. *Artificial Intelligence* 24(1-3): 85–168. doi.org/10.1016/0004-3702(84)90038-9
- Forbus, K. 2001. Exploring Analogy in the Large. In *Analogy: Perspectives from Cognitive Science*, ed. D. Gentner, K. Holyoak, and B. Kokinov. Cambridge, MA: The MIT Press.
- Forbus, K., 2014. Qualitative Reasoning. In *Computing Handbook*, 3rd Edition: Computer Science and Software Engineering. Boca Raton, FL: CRC Press. doi.org/10.1201/b16812-41
- Forbus, K. 2015. Analogical Abduction and Prediction: Their Impact on Deception. In *Deceptive and Counter-Deceptive Machines: Papers from the AAAI Fall Symposium*. Technical Report FS-15-03. Palo Alto, CA: AAAI Press.
- Forbus, K. 2016. Software Social Organisms: Implications for Measuring AI Progress. *AI Magazine* 37(1): 85–90. doi.org/10.1609/aimag.v37i1.2648
- Forbus, K., and de Kleer, J. 1994. *Building Problem Solvers*. Cambridge, MA: The MIT Press.
- Forbus, K.; Nielsen, P.; and Faltings, B. 1991. Qualitative Spatial Reasoning: The CLOCK Project. *Artificial Intelligence* 51 (1-3): 417–471. doi.org/10.1016/0004-3702(91)90116-2
- Forbus, K.; Riesbeck, C.; Birnbaum, L.; Livingston, K.; Sharma, A.; and Ureel, L. 2007. Integrating Natural Language, Knowledge Representation and Reasoning, and Analogical Processing to Learn by Reading. In *Proceedings of the Twenty-Second Conference on Artificial Intelligence*. Palo Alto, CA: AAAI Press.
- Forbus, K.; Ferguson, R.; Lovett, A.; and Gentner, D. 2016. Extending SME to Handle Large-Scale Cognitive Modeling. *Cognitive Science* 41(5): 1152–1201. doi.org/10.1111/cogs.12377
- Forbus, K. and Gentner, D. 1997. Qualitative Mental Models: Simulations Or Memories? Paper presented at the Eleventh International Workshop on Qualitative Reasoning, Cortona, Italy, June 3-6.
- Forbus, K.; Gentner, D.; and Law, K. 1995. MAC/FAC: A Model of Similarity-Based Retrieval. *Cognitive Science* 19(2)(April–June): 141–205. doi.org/10.1207/s15516709cog1902_1
- Forbus, K.; Klenk, M.; and Hinrichs, T. 2009. Companion Cognitive Systems: Design Goals and Lessons Learned So Far. *IEEE Intelligent Systems* 24(4): 36–46. doi.org/10.1109/MIS.2009.71
- Forbus, K.; Liang, C.; and Rabkina, I. 2017. Representation and Computation in Cognitive Models. *Topics in Cognitive Science*. doi.org/10.1111/tops.12277
- Forbus, K.; Usher, J.; Lovett, A.; Lockwood, K.; and Wetzel, J. 2011. CogSketch: Sketch Understanding for Cognitive Science Research and for Education. *Topics in Cognitive Science* 3(4): 648–666. doi.org/10.1111/j.1756-8765.2011.01149.x
- Friedman, S. E. and Forbus, K. 2010. An Integrated Systems Approach to Explanation-Based Conceptual Change. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*. Palo Alto, CA: AAAI Press.
- Friedman, S., and Forbus, K. 2011. Repairing Incorrect Knowledge with Model Formulation and Metareasoning. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, Barcelona, Spain. Palo Alto, CA: AAAI Press.
- Friedman, S.; Barbella, D.; and Forbus, K. 2012. Revising Domain Knowledge with Cross-Domain Analogy. *Advances in Cognitive Systems* 2: 13–24.
- Friedman, S. E.; Forbus, K. D.; and Sherin, B. 2011. How Do the Seasons Change? Creating and Revising Explanations via Model Formulation and Metareasoning. Paper presented at the 25th International Workshop on Qualitative Reasoning, Barcelona, Spain, 16–18 July.
- Gentner, D. 1983. Structure-Mapping: A Theoretical Framework for Analogy. *Cognitive Science* 7(2): 155–170. doi.org/10.1207/s15516709cog0702_3
- Gentner, D.; Brem, S.; Ferguson, R. W.; Markman, A. B.; Levidow, B. B.; Wolff, P.; and Forbus, K. D. 1997. Analogical Reasoning and Conceptual Change: A Case Study of Johannes Kepler. *The Journal of the Learning Sciences* 6(1): 3–40. doi.org/10.1207/s15327809jls0601_2
- Goldstone, R. L.; Medin, D. L.; and Gentner, D. 1991. Relational Similarity and the Non-independence of Features in Similarity Judgments. *Cognitive Psychology* 23(2): 222–264. doi.org/10.1016/0010-0285(91)90010-L
- Hinrichs, T., and Forbus, K. 2011. Transfer Learning Through Analogy in Games. *AI Magazine* 32(1): 72–83. doi.org/10.1609/aimag.v32i1.2332
- Hinrichs, T., and Forbus, K. 2012a. Learning

- Qualitative Models by Demonstration. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence*, 207–213. Palo Alto, CA: AAAI Press.
- Hinrichs, T., and Forbus, K. 2012b. Toward Higher-Order Qualitative Representations. Paper presented at the Twenty-Sixth International Workshop on Qualitative Reasoning, Los Angeles CA, July 16–18.
- Hinrichs, T., and Forbus, K. 2014. X goes First: Teaching Simple Games Through Multimodal Interaction. *Advances in Cognitive Systems* 3: 31–46.
- Hinrichs, T., and Forbus, K. 2016. Qualitative Models for Strategic Planning. *Advances in Cognitive Systems* 4: 75–92.
- Ioannides, C., and Vosniadou, S. 2002. The Changing Meanings of Force. *Cognitive Science Quarterly* 2: 5–61.
- Kahneman, D. 2011. *Thinking, Fast and Slow*. London: Allen Lane.
- Kandaswamy, S.; Forbus, K.; and Gentner, D. 2014. Modeling Learning via Progressive Alignment Using Interim Generalizations. In *Proceedings of the 36th Annual Meeting of the Cognitive Science Society*, Quebec City, Quebec, Canada, 23–26 July. Austin, TX: Cognitive Science Society Inc.
- Klenk, M., and Forbus, K. 2009. Analogical Model Formulation for AP Physics Problems. *Artificial Intelligence*, 173(18), 1615–1638. doi.org/10.1016/j.artint.2009.09.003
- Klenk, M., and Forbus, K. 2013. Exploiting Persistent Mappings in Cross-Domain Analogical Learning of Physical Domains. *Artificial Intelligence* 195: 398–417. doi.org/10.1016/j.artint.2012.11.002
- Kuehne, S.; Gentner, D.; and Forbus, K. 2000. Modeling Infant Learning via Symbolic Structural Alignment. In *Proceedings of the 22nd Annual Meeting of the Cognitive Science Society*, Philadelphia, PA, August. Austin, TX: Cognitive Science Society Inc.
- Laird, J. 2012. *The SOAR Cognitive Architecture*. Cambridge, MA: The MIT Press
- Laird, J.; Lebiere, C.; and Rosenbloom, P. 2017. A Standard Model of the Mind: Toward a Common Computational Framework Across Artificial Intelligence, Cognitive Science, Neuroscience, and Robotics. *AI Magazine* 38(4). doi.org/10.1609/aimag.v38i4.2744
- Liang, C., and Forbus, K. 2015. Learning Plausible Inferences from Semantic Web Knowledge by Combining Analogical Generalization with Structured Logistic Regression. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*. Palo Alto, CA: AAAI Press.
- Liang, C., Paritosh, P., Rajendran, V., and Forbus, K. D. 2016. Learning Paraphrase Identification with Structural Alignment. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, New York, NY USA. Palo Alto, CA: AAAI Press.
- Lockwood, K., and Forbus, K. 2009. Multimodal Knowledge Capture from Text and Diagrams. In *Proceedings of the Fifth International Conference on Knowledge Capture (KCAP-2009)*. New York: Association for Computing Machinery. doi.org/10.1145/1597735.1597747
- Lovett, A., and Forbus, K. 2017. Modeling Visual Problem Solving as Analogical Reasoning. *Psychological Review* 124(1): 60–90. doi.org/10.1037/rev0000039
- Lovett, A., and Forbus, K. 2011. Cultural Commonalities and Differences in Spatial Problem-Solving: A Computational Analysis. *Cognition* 121(2): 281–287. doi.org/10.1016/j.cognition.2011.06.012
- Lovett, A.; Tomai, E.; Forbus, K.; and Usher, J. 2009b. Solving Geometric Analogy Problems Through Two-Stage Analogical Mapping. *Cognitive Science* 33(7): 1192–1231. doi.org/10.1111/j.1551-6709.2009.01052.x
- Markman, A. B., and Gentner, D. 1993. Structural Alignment during Similarity Comparisons. *Cognitive Psychology* 25(4): 431–467. doi.org/10.1006/cogp.1993.1011
- Markman, A. B., and Gentner, D. 1996. Commonalities and Differences in Similarity Comparisons. *Memory and Cognition* 24(2): 235–249. doi.org/10.3758/BF03200884
- Markman, A. B., and Medin, D. L. 2002. Decision Making. In *Stevens' Handbook of Experimental Psychology* 3rd edition. New York: John Wiley & Sons, Inc. doi.org/10.1002/0471214426.pas0210
- McFate, C., and Forbus, K. 2016a. Analogical Generalization and Retrieval for Denominal Verb Interpretation. In *Proceedings of the 38th Annual Meeting of the Cognitive Science Society*, Philadelphia, PA, 10–13 August. Austin, TX: Cognitive Science Society, Inc.
- McFate, C., and Forbus, K. 2016b. An Analysis of Frame Semantics of Continuous Processes. In *Proceedings of the 38th Annual Meeting of the Cognitive Science Society*, Philadelphia, PA, 10–13 August. Austin, TX: Cognitive Science Society, Inc.
- McFate, C. J.; Forbus, K.; and Hinrichs, T. 2014. Using Narrative Function to Extract Qualitative Information from Natural Language Texts. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*. Palo Alto, CA: AAAI Press.
- McLure, M. D.; Friedman S. E.; and Forbus, K. D. 2015. Extending Analogical Generalization with Near-Misses. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*. Palo Alto, CA: AAAI Press.
- McShane, M. 2017. Natural Language Understanding (NLU, not NLP) in Cognitive Systems. *AI Magazine* 38(4). doi.org/10.1609/aimag.v38i4.2745
- Mitchell, T.; Cohen, W.; Hruschka, E.; Talukdar, P.; Betteridge, J.; Carlson, A.; Dalvi, B.; Gardner, M.; Kisiel, B.; Krishnamurthy, J.; Lao, N.; Mazaitis, K.; Mohamed, T.; Nakashole, N.; Plataniotis, E.; Ritter, A.; Samadi, M.; Settles, B.; Wang, R.; Wijaya, D.; Gupta, A.; Chen, X.; Saparov, A.; Greaves, M.; Welling, J. 2015. Never-Ending Learning. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*. Palo Alto, CA: AAAI Press.
- Newell, A. 1990. *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.
- Sagi, E.; Gentner, D.; and Lovett, A. 2012. What Difference Reveals About Similarity. *Cognitive Science* 36(6): 1019–1050. doi.org/10.1111/j.1551-6709.2012.01250.x
- Scheutz, M. 2017. The Case for Explicit Ethical Agents. *AI Magazine* 38(4). doi.org/10.1609/aimag.v38i4.2746 s
- Tomai, E., and Forbus, K. 2007. Plenty of Blame to Go Around: A Qualitative Approach to Attribution of Moral Responsibility. Paper presented at the Twenty-First Workshop on Qualitative Reasoning, Aberystwyth, UK, 27–29 June.
- Wilson, J.; Forbus, K.; and McLure, M. 2013. Am I Really Scared? A Multiphase Computational Model of Emotions. In *Proceedings of the Second Conference on Advances in Cognitive Systems*, Baltimore, MD, 12–14 December. Austin, TX: Cognitive Science Society Inc.
- Winston, P. H. 1970. Learning Structural Descriptions From Examples. Ph.D. diss., Computer Science, Massachusetts Institute of Technology, Cambridge, MA.

Kenneth D. Forbus is the Walter P. Murphy Professor of Computer Science and a professor of education at Northwestern University. His research interests include qualitative reasoning, analogical reasoning and learning, natural language understanding, sketch understanding, and cognitive architecture.

Thomas Hinrichs is a research associate professor of computer science at Northwestern University. His research interests include commonsense reasoning, cognitive architectures and machine learning.