

Learning Causal Models via Progressive Alignment & Qualitative Modeling: A Simulation

Scott E. Friedman (friedman@northwestern.edu)

Qualitative Reasoning Group, Northwestern University, 2145 Sheridan Rd
Evanston, IL 60208-0834 USA

Kenneth D. Forbus (forbus@northwestern.edu)

Qualitative Reasoning Group, Northwestern University, 2145 Sheridan Rd
Evanston, IL 60208-0834 USA

Abstract

Learning causal models is one of the central problems of cognitive science. We describe a simulation of early learning in physical domains from observations that uses progressive alignment and qualitative modeling to derive plausible causal models from observations. We show how *protohistories* can be created via progressive alignment and used with covariance algorithms to infer causality. The result, a *causal corpus* described using qualitative representations, can make simple predictions and set the stage for more sophisticated later models. The simulation has been successfully tested with three learning problems, with encouraging results.

Keywords: Cognitive modeling, causal reasoning, domain learning, qualitative modeling

Introduction

Learning causal models is a central problem in cognitive science. We are building on the framework of Forbus & Gentner (1986), which proposed decomposing learning of physical domains from experience into four stages. (1) *Protohistories* are prototypical behaviors, generalized from multiple experiences. (2) The *causal corpus* consists of fragmentary causal models, created from protohistories. (3) These fragmentary models are organized into a *naïve physics*, which regularizes the fragmentary causal models by postulating broadly applicable mechanisms. (4) *Expert understanding* consists of deepening the naïve physics and tying it to mathematical and other formal models, typically culturally learned. Importantly, these stages are localized, within the understanding of particular phenomena (compatible with notions of situated cognition, e.g. Brown *et al* (1989)). For example, someone might have an expert understanding of electronics while having only a partial set of causal models for fluids.

In this paper, we focus on learning plausible causal models from observations, stages (1) and (2) above. We are using ideas from qualitative process theory (Forbus, 1984) to formally represent partial causal models. We utilize analogical learning techniques from structure-mapping theory (Gentner, 1983), particularly progressive alignment (Kotovsky & Gentner, 1996), plus statistical methods to construct the partial causal models from representations of experiences. The simulation has been successfully tested on learning problems from three domains; we use understanding floating versus sinking as a running example

for illustration. We first review QP theory, the structure-mapping models our simulation uses, and our use of covariance in making causal hypotheses. Then we discuss how protohistories are learned from experience via progressive alignment, proposing the idea of *generalization contexts* as a means of organizing experience around salient questions. Next we discuss three strategies for constructing fragmentary causal models, which hypothesize conditions under which prototypical behaviors occur, limit points, and causal relationships between quantities. We summarize results from the simulation. Finally, we discuss other related work and future plans.

Background

Our theoretical framework uses *qualitative process theory* (Forbus, 1984) as its account of causality. QP theory concerns natural changes that occur in dynamic systems, independent of agents. Changes are caused by *physical processes*, which provide the notion of mechanism for causality (cf. Chi *et al* 1994; Ahn *et al* 1995)¹. Heat flow and boiling are two examples of physical processes. Processes directly cause changes in some quantities, e.g. heat flow directly influences the heat of its source and its destination. These changes propagate through the system via *qualitative proportionalities* which are causal relationships between quantities. For instance, since *density* α_{Q+} *mass* (density is qualitatively proportional to mass) is true, causing the mass of a gas to increase will cause its density to increase, assuming all else remains constant. Qualitative proportionalities provide only partial information about what will happen, as opposed to more monolithic rules. This makes them particularly appropriate for representing local causal models, since models learned from one set of experiences can be more easily combined with others.

These causal laws are contextualized by belonging to either processes or *views*. Views characterize states of entities and constellations of entities. Both processes and views hold only when their *conditions* are true. An important kind of condition is ordinal relations between

¹ QP theory has been used in causal modeling in non-physical domains, like economics and political reasoning, hence *continuous processes* might be more appropriate.

quantities. For example, water boils when its temperature reaches its boiling point. Parameters like boiling points, which determine when processes start/stop and views hold or not are called *limit points*. Postulating the existence of limit points is an important challenge in learning QP models, since they are crucial for prediction.

QP theory provides an account of how causal models are represented and reasoned with, but it does not describe how they are learned. We claim that statistical accounts of causality (cf. Pearl, 2000; Glymour & Cheng, 1998; Gopnik *et al* 2004) provide techniques that can be harnessed to produce QP models. We incorporate statistics via similarity, using structure-mapping operations to construct probabilities as a side-effect of assimilating experiences. The SEQL model of generalization (Kuehne *et al* 2000) constructs generalizations incrementally via analogical comparison. We simulate analogical matching via SME, the Structure-Mapping Engine (Falkenhainer *et al* 1989; Forbus *et al* 1994). Given two structured representations, the *base* and *target*, SME computes one or two *mappings* which describe how the base and target can be aligned. Mappings include a set of *correspondences* that detail exactly which entities and statements in one description go with entities and statements in the other, a *structural evaluation score* which indicates the overall quality of the match, and a set of *candidate inferences* that are conjectures about the target, using the correspondences to project partially unmapped base structures. Candidate inferences allow predictions and explanations to be generated without rules, via analogy to prior experiences and explanations. This makes them particularly important for accounts of learning like ours that postulate localized, incrementally generated models.

SEQL operates by maintaining a list of generalizations and exemplars. Given a new exemplar, SEQL compares it with the generalizations. If it is sufficiently similar to one of them, it is assimilated into that generalization. If the new exemplar doesn't fit an existing generalization, it is compared against the list of unassimilated exemplars. If a pair of exemplars is sufficiently similar, they are combined to form a new generalization. The combination and assimilation processes are basically the same: SEQL maintains a probability for each statement in the generalization, which is updated as each new exemplar is assimilated (Halstead & Forbus, 2005). These probabilities are exploited in our simulation to do statistical learning.

We call a set of generalizations and exemplars that are being processed together by SEQL a *generalization context*. Generalization contexts can be defined bottom-up, via similarity-based retrieval, or by labeling, e.g., a learner might use a generalization context to process all of the examples that have been given a verbal label, like "cat".

Learning Protohistories

Protohistories are generalizations of specific observed behaviors. Observed behaviors are typically rich with perceptual information, and in new domains, impoverished

with regard to explanations. We postulate that analogical generalization, as modeled in SEQL, is used to construct prototypical behaviors. Here is an example of an observation given to our simulation. It describes an adult female human, swimming (gliding) in a still pond, and floating:

```
(isa bodyInLiquid0 AdultFemaleHuman)
(isa container0 Pond)
(isa liquid0 (LiquidFn Water))
(in-UnderspecifiedContainer liquid0 container0)
(massOfObject bodyInLiquid0 (Kilogram 60))
(volumeOfObject bodyInLiquid0 (CubicCentimeter 62039))
(isa gliding0 MovementEvent)
(primaryObjectMoving gliding0 bodyInLiquid0)
(isa stillLiquid0 StandingStill)
(doneBy stillLiquid0 liquid0)
(in-Floating bodyInLiquid0 liquid0).
```

The vocabulary of concepts and relations is drawn from the ResearchCyc knowledge base², an independently developed representation system for common-sense knowledge. The predicate calculus was produced using a natural-language understanding (NLU) system from simplified English (Kuehne & Forbus, 2004), to reduce tailorability.

The simplified English that generates the above predicate calculus observation is:

The woman *bodyInLiquid0* floats in water *liquid0* in a pond *container0*. The mass of the woman *bodyInLiquid0* is 60 kilograms. The volume of the woman *bodyInLiquid0* is 62039 cubic centimeters. The woman *bodyInLiquid0* is moving but the water *liquid0* is standing still.

The NLU system first translates this English into a set of semantic choices for manual disambiguation. Next, the system generates a predicate calculus interpretation, per the above example input. In the above passage, the italicized words are used to maintain referential continuity across sentences; the NLU system utilizes these words as entity names. Some of the entities (such as *stillLiquid0* and *gliding0*) were manually translated for illustration.

For SME processing, *isa* statements are automatically translated into attributes (i.e., (*AdultFemaleHuman bodyInLiquid0*)). SEQL generalizations abstract specific individuals (e.g., *bodyInLiquid0*) into anonymous individuals, not variables. Numerical parameters (e.g., (*Kilogram 60*)) are also abstracted into anonymous individuals, but their values are preserved in a distribution for each quantity in the generalization. These distributions are used to conjecture limit points below. We ignore memory retrieval in this simulation, and provide as input a stream of observations like the above.

How many generalization contexts should be used? Since SEQL automatically constructs multiple generalizations according to similarity, one possibility is to use a single context. The drawback with a single context is that it does

² www.cycorp.com

not necessarily provide the discrimination one needs for learning. For example, situations involving boats, people, and leaves can include both cases of floating and cases of sinking. We have observed that SEQL will, because of attribute information, cluster such situations by the kinds of entities involved. (This tendency would be even stronger with more perceptual information in the representations.) A protohistory that includes both floating and sinking cases will not be useful for learning conditions for predicting which will occur. Consequently, we assume that when a learner wants to understand how to discriminate between two salient possibilities, such as whether something floats or sinks when placed in water, that they create separate generalization contexts for each possibility. Each such generalization context incorporates a set of *entry patterns* that are tested against new exemplars. When a new exemplar satisfies the entry pattern for a generalization context, it is processed in that context. These entry patterns allow exclusivity of exemplars that might otherwise be generalized together. The same exemplar can be processed in multiple contexts, since a learner might be trying to understand multiple concepts at once.

Consider a learner trying to understand the distinction between floating and sinking, as well as sailboats sailing. Figure 1 illustrates the three example generalization contexts that would be used. If an exemplar arrives with (SinkingEvent sinking0) as a constituent fact, with no mention of floating, it will be incorporated into the rightmost context alone. If another exemplar arrives with (isa boat0 SailBoat) and (floating-in boat0 (LiquidFn Water)) as constituent facts, it will be incorporated into both leftmost and middle contexts.




Figure 1: Example contextual protohistory organization

Learning a Causal Corpus

The causal corpus consists of a set of causal models grounded in, and connecting, protohistories. These causal models are local to particular protohistories or collections of protohistories – restructuring these local models into general domain theories, as per Forbus & Gentner (1986), occurs only after a reasonable causal corpus has been constructed.

Even fragmentary causal models are quite powerful: Understanding what qualitative proportionalities hold in a protohistory yields a means of predicting the immediate consequences of parameter changes. Similarly, understanding quantity conditions that determine which protohistory represents the behavior that occurs in a situation enables predictions of state changes.

Our simulation uses a set of four *causal learning strategies* – procedures that take protohistories and quantities as input, and annotate those protohistories with causal hypotheses, expressible using the vocabulary of QP theory. We also describe a strategy for deriving complex quantities from constituent input quantities. We do not view this set of strategies as complete, but we believe they are a good starting point.

Unary Causal Learning Strategy

The *unary causal learning strategy* postulates a single property or relationship as the cause of a phenomenon. An example is the naïve statement *the entity sinks because it is heavy*, which we can express as:

```
(CAUSE
  (hasAttribute entity1 Heavy)
  (AND
    (SinkingEvent sinking0)
    (doneBy sinking0 entity1)))
```

This strategy is based on the covariance of expressions across protohistories, where high-magnitude covariance implies a belief that a causal mechanism exists. We compute covariance using the ΔP algorithm (Lopez *et al*, 1998). For cause c and effect e , let $P(e^+|c^+)$ be the probability of the effect given the cause is present, and $P(e^+|c^-)$ be the probability of the effect given the cause is absent.

$$\Delta P = P(e^+|c^+) - P(e^+|c^-).$$

ΔP is the probability of the effect given the cause, less the probability of the effect given the absence of the cause. We calculate probabilities with the following equation which normalizes by $|i|$, the number of exemplars in the protohistory i :

$$P(e) = \sum P_i(e) * |i| / \sum |i|.$$

To filter out low-magnitude covariance, we employ a cutoff-value k , $1 > k > 0$, such that if $\Delta P > k$, c is a generative cause and if $\Delta P < -k$, c is a preventative cause.

The unary causal strategy may produce contradictory causal stories and may be used to simulate misconceptions, as we demonstrate in our results below.

Quantity Analysis

Analyzing quantity values enables us to hypothesize limit points, quantity conditions, and qualitative proportionalities. The *quantity condition strategy*

identifies relevant ordinal relationships. The *limit point strategy* hypothesizes new causally-relevant values. The *quantity derivation strategy* hypothesizes compound quantities. We discuss each in turn.

Quantity Condition Strategy. Conditions for processes and views typically include ordinal relations between quantities. For instance, for a body to be floating in a liquid, its density must be less than the liquid’s density. Quantity conditions are conjectured as follows:

1. Protohistories that summarize experience related to the target phenomenon are divided into two groups: those that express the phenomenon (P^+) and those that do not (P^-).

2. For each protohistory p_i within ($P^+ \cup P^-$), the ordinal relationships $R_i = \{r_1, r_2, \dots, r_n\}$ are identified that hold for every exemplar within P_i . The ordinal relationships tested are $=, >, <, \geq,$ and $\leq,$ over the set of exemplars that were used in forming P_i .

3. Conditions are identified that pertain to the entirety of P^+ and P^- , such that $R^+ = \{R^+_1 \cap \dots \cap R^+_n\}$ and $R^- = \{R^-_1 \cap \dots \cap R^-_n\}$.

4. Conditions that coincide with the phenomenon are the set $R_{\text{cause}} = R^+ - R^-$. Relationships that coincide with the absence of the phenomenon are the set $R_{\text{prevent}} = R^- - R^+$.

We resort to exemplars in step 2 because our encoding process does not automatically generate ordinal relationships from numerical values in observations. (The quantity value distribution information stored with generalizations cannot be used to compute this, because links to particular exemplars is not part of that distribution.) This is a simplification: We believe that psychologically, encoding choices are driven in part by learning goals, which would propose encoding particular ordinal relationships in order to test conjectures via this strategy. Such goals might be generated based on trying various ordinals on a small number of exemplars, but that is left for future work.

Limit Point Strategy. Some physical phenomena occur when a quantity’s value is above or below a specific limit point. Like the quantity condition strategy, the limit point strategy assumes that two sets of protohistories have been identified, such as water being heated and boiling and water being heated and not boiling. Recall that protohistories preserve the set of exemplar values $\{v_1, v_2, \dots, v_n\}$ for each quantity. This information can be summarized via an interval V , where $V = [\min(v_1, v_2, \dots, v_n), \max(v_1, v_2, \dots, v_n)]$.

After calculating quantity intervals for individual protohistories, we first compute possible limit points by grouping protohistories into two sets: those that express the given phenomena $P^+ = \{p^+_1, p^+_2, \dots, p^+_n\}$ and those that do not $P^- = \{p^-_1, p^-_2, \dots, p^-_n\}$. For each quantity-type q , we merge the protohistory intervals so that

$$P^+_q = [\min(p^+_1q, p^+_2q, \dots, p^+_nq), \max(p^+_1q, p^+_2q, \dots, p^+_nq)]$$

$$P^-_q = [\min(p^-_1q, p^-_2q, \dots, p^-_nq), \max(p^-_1q, p^-_2q, \dots, p^-_nq)].$$

If the intervals P^+_q and P^-_q do not overlap for a quantity, it could be the case that a limit point exists within the interval $[\max(\min(P^+_q, P^-_q)), \min(\max(P^+_q, P^-_q))]$, or between the maximum point of the lower interval and the minimum point of the higher interval. This interval is then added to the causal corpus, as a limit point approximation.

If the intervals P^+_q and P^-_q overlap, there could still be an uninterrupted interval $[q_{\min}, q_{\max}]$ that represents a condition under which the physical phenomenon occurs. Instead of merging protohistory intervals into P^+_q and P^-_q , we test for exclusiveness, such that no protohistory intervals in P^+ overlap protohistory intervals in P^- for a quantity q . Uninterrupted intervals in q are then added to the causal corpus as possible conditions for the target phenomenon.

Quantity Derivation Strategy. Understanding many physical phenomena requires introducing quantities beyond those observed. To understand why something floats versus sinks, for example, requires introducing the idea of density. If the quantity analysis fails to distinguish between two behaviors within the observable quantities, the quantity derivation strategy proposes new quantities which are then searched for limit points and ordinal relationships. For all explicitly mentioned quantities a and b such that $a \neq b$, a set of new quantities C is derived:

$$C = \{a/b, b/a, a*b, a+b, a-b, b-a\}.$$

The units for the derived quantities may be identical to their constituent quantities ($kg + kg = kg$), or they may be combinations of their constituent units ($kg/cc = kg/cc$).

Simulation Results

Now we show how these methods combine to produce plausible causal corpus elements from a set of observations. We first go through a single learning task in detail, then summarize the results of others.

To investigate learning the floating versus sinking distinction, we encoded 30 unique exemplars – 16 floating and 14 sinking – in simplified English, which was fed into a natural language understanding system to automatically produce predicate calculus descriptions like our earlier example. Many factors used in the scenarios were based on Piaget’s (1930) interviews with children: the motion of the water (still or wavy); the body in water (man, woman, log, cruise ship, or tree branch); the body of water (ocean, sea, lake, pond, bath-tub, or bowl); and autonomous motion of the body (moving/gliding or still). In the scenarios, a body floats when: (1) the body moves autonomously; (2) the body of water has waves; and (3) the body’s density is less than 1 g/cc. Phenomena (1) and (2) are misconceptions reported in Piaget’s (1930) interviews with children; phenomenon (3) corresponds to physical law.

The simulation first generates protohistories from the exemplars. Two generalization contexts were used, with entry patterns `(in-Floating ?x ?y)` and `(isa ?x SinkingEvent)`, to model the focus on understanding when

something floated or sank. The assimilation threshold for SEQL was set to 0.75. This yielded six protohistories, five for floating and one for sinking. All exemplars were assimilated into a generalization. Table 1 shows the protohistory abstractions with the generic entities in bold, and the protohistory size, |P|. Protohistories P₁ and P₄ preserved *tree branch* and *cruise ship* in their abstractions, respectively; the rest of the protohistories contain only generic entities.

Context	#	Protohistory Abstraction	P
Floating	1	Idle tree branch, wavy water	2
	2	Moving body	3
	3	Moving body	3
	4	Moving cruise ship	3
	5	Wavy water	5
Sinking	6	Idle body , still water	14

Table 1: Protohistories for floating and sinking

The abstractions for P₂ and P₃ are identical, yet they are still different protohistories. This is due to uncertain facts within the generalization. Specifically, in P₂, P(body = man) = .66, and in P₃, P(body = woman) = .66. Thus, although the abstractions are identical, the underlying representations differ. Low-probability facts are considered for similarity processing, so they remained distinct.

To generate causal corpus information for these protohistories, the strategies defined above were executed in the order given. Table 2 shows the causal hypotheses generated by the unary strategy, where positive covariation indicates that the factor enables floating, while negative covariation indicates that the factor prevents it.

Phenomenon	Covariance with Floating (ΔP)
Autonomously moving body	0.67
Still body	-0.88
Presence of waves	0.5
Still water	-1.0
Body = Cruise Ship	0.54
Body = Tree Branch	0.5

Table 2: Covariance results for floating.

As expected, these misconceptions correspond to those found in Piaget’s (1930) interviews with children. Given that the stimuli were also designed based on this interview data, this is not surprising, but it is reassuring given that (a) the stimuli themselves were generated semi-automatically and (b) there were distractors present.

Given a set of protohistories, the simulation proceeds to analyze its quantities, searching for limit points and quantity conditions that help explain floating. The observable quantities yielded no causal hypotheses, so the simulation used the quantity derivation strategy to create new quantities and try again. One of the derived quantities does yield a limit point, as shown in Table 3. Since this limit point

(which we know as density) was derived as the ratio of mass and volume, we also obtain the qualitative proportionalities shown in Table 3, imposing a causal direction on what was an algebraic relationship by assuming that observable parameters are more primitive than derived parameters. (This is a heuristic, of course, that could be incorrect – consider heat derived from temperature, for example.)

Type of Causal Hypothesis	Formula
Derived Quantity	$q = \text{mass}_{\text{body}} / \text{volume}_{\text{body}}$
Limit Point	$q < [0.001, 0.00102] \text{ kg/cc}$
Qualitative Proportionality	$\text{floatability} \propto_{Q-} q$
	$q \propto_{Q+} \text{mass}_{\text{body}}$
	$q \propto_{Q-} \text{volume}_{\text{body}}$
	$\text{floatability} \propto_{Q-} \text{mass}_{\text{body}}$
	$\text{floatability} \propto_{Q+} \text{volume}_{\text{body}}$

Table 3: Causal hypotheses generated about floating

In addition to floating/sinking, we have tested the simulation on two other learning scenarios. To model learning how balance scales work (Siegler, 1983), we encoded nine scenarios using the methodology above, varying the kinds of objects on the balance and the posture of the object (e.g., sitting or kneeling or upright). Using two generalization contexts, one for right-side sinking and one for left-side sinking, the simulation generated two protohistories for each context. The unary strategy produced an interesting misconception:

Delta-P(0.71): (postureOfAnimal rightobj0 UprightPosture).

The quantity condition strategy creates the sensible quantity hypothesis

(> (massOfObject leftside0) (massOfObject rightside0))

to predict when the left side will sink.

In another learning experiment conjecturing when boiling would occur, six exemplars were encoded using the methodology above. The limit point strategy conjectures a limit point for temperature to predict when boiling occurs:

Hypothesis: phenomena occurs when
(temperatureOfObject kettle0)
is above some point in the range:
[95.0-100.0] DegreeCelsius.

The lower bound of the range could be refined by more experience, of course.

While the number of learning experiments conducted to date is small, the results obtained so far are very reasonable.

Related Work

Some of diSessa’s (1983) p-prims (for “phenomenological primitives”) can be viewed as causal corpus elements while others may be viewed as protohistories. No computational model for learning them was ever implemented.

The closest previous simulations are COBWEB system (Fisher, 1987) which utilized conceptual clustering and INTHELEX (Esposito *et al*, 2000), which revised theories described as prolog programs. Neither simulation introduced causal models, nor was tested on semi-automatically generated stimuli. Our quantity derivation strategy is inspired by Langley's (1981) BACON simulation.

The process of inferring causality has been modeled with statistical methods such as Bayesian reasoning (Schulz, Bonawitz, & Griffiths, 2007; Gopnik *et al*, 2004) and causal strength (Lopez *et al*, 1998). Our covariance algorithm exploits such techniques, but we believe our use of QP theory to express the results of causal hypotheses is more psychologically plausible.

Discussion and Future Work

Our simulation combines symbolic, relational causal models with statistical inference to learn them. We think this is a very promising approach to modeling learning more generally. To be sure, many advocates of statistical accounts of causality do not include any notion of mechanism, and we obviously (along with Chi *et al* 1994; Ahn *et al* 1995) do not believe that is sufficient. Moreover, there are suggestions that people may infer correlations and probabilities from their knowledge of causal mechanisms, rather than the other way around (Rips, in press; Tversky & Kahneman, 1980). Identifying how each contributes to causal reasoning and learning appears to be a very promising approach.

However, this simulation is obviously only a beginning. In addition to testing the simulation on a broader range of learning problems, we also plan to incorporate retrieval, building on MAC/FAC (Forbus *et al* 1995). Having the simulation generate its own distinctions to explore, perhaps via failed predictions made with protohistories, is also an important problem to explore.

References

- Ahn, W. k., Kalish, C. W., Medin, D. L., & Gelman, S. A. (1995). The role of covariation versus mechanism information in causal attribution. *Cognition*, *54*, 299-352.
- Brown, J.S., Collins, A., & Duguid, S. (1989). Situated cognition and the culture of learning. *Educational Researcher*, *18*(1), 32-42.
- Chi, M., Slotta, J., & De Leeuw, N. (1994). From things to processes: A theory of conceptual change for learning science concepts. *Learning and Instruction*, *4*(1), 27-43.
- diSessa, A. (1983). Phenomenology and the evolution of intuition. In D. Gentner and A. Stevens (Eds), *Mental Models*. 15-33. Hillsdale, NJ: Lawrence Erlbaum.
- Esposito, F., Semeraro, G., Fanizzi, N., & Ferilli, S. (2000). Conceptual Change in Learning Naive Physics: The Computational Model as a Theory Revision Process. In E. Lamma and P. Mello (Eds.), *AI*IA99: Advances in Artificial Intelligence, Lecture Notes in Artificial Intelligence 1792*, 214-225, Springer: Berlin.
- Fisher, D. H. (1987). Knowledge acquisition via incremental concept clustering. *Machine Learning*, *2*, 139-172.
- Forbus, K. (1984). Qualitative Process Theory. *Artificial Intelligence*, *24*, 85-168.
- Forbus, K. D., & Gentner, D. (1986). Learning physical domains: towards a theoretical framework. In R.S. Michalski, J. G. Carbonell and T. M. Mitchell *Machine Learning: An Artificial Intelligence Approach, Vol. 2*.
- Forbus, K., Gentner, D. & Law, K. (1995). MAC/FAC: A model of Similarity-based Retrieval. *Cognitive Science*, *19*(2), 141-205.
- Glymour, C., & Cheng, P. (1998). Causal mechanism and probability: a normative approach. In: M. Oaksford and N. Chater, Editors, *Rational Models of Cognition*, Oxford University Press.
- Gopnik, A., Glymour, C., Sobel, D., Schulz, L., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, *111*, 1, 1-31.
- Halstead, D. & Forbus, K. (2005). Transforming between Propositions and Features: Bridging the Gap. Proceedings of AAAI-2005. Pittsburgh, PA.
- Kotovskiy, L., & Gentner, D. (1996). Comparison and categorization in the development of relational similarity. *Child Development*, *67*, 2797-2822.
- Kuehne, S. & Forbus, K. (2004). Capturing QP-relevant information from natural language text. *Proceedings of the 18th International Qualitative Reasoning Workshop*.
- Kuehne, S. E., Forbus, K. D., Gentner, D., & Quinn, B. (2000). SQL: Category learning as progressive abstraction using structure mapping. *Proceedings of CogSci 2000*.
- Langley, P. (1981). Data-driven discovery of physical laws. *Cognitive Science* *5*, 31-54.
- Lopez, F. J., Cobos, P. L., Cano, A., & Shanks, D. R. (1998). The rational analysis of human causal and probability judgement. In M. Oaksford, N. Chater *et al Rational Models of Cognition*, Oxford University Press.
- Pearl, J. (2000). *Causality: Models, Reasoning, and Inference*. Cambridge University Press.
- Piaget, J. (1930). *The Child's Conception of Physical Causality*. Kegan Paul, London.
- Rips, L. J. (in press). Causal thinking. In J. E. Adler & L. J. Rips (Eds.), *Reasoning: Studies of Human Inference and Its Foundation*. Cambridge, UK: Cambridge University Press.
- Schulz, L. E., Bonawitz, E. B., & Griffiths, T. L. (2007). Can being scared make your tummy ache? Naive theories, ambiguous evidence, and preschoolers' causal inferences. *Developmental Psychology*, *43*, 1124-1139.
- Siegler, R. (1983). Five generalizations about cognitive development. *American Psychologist*, *38*, 263-277.
- Tversky, A., & Kahneman, D. (1980). Causal schemas in judgment under uncertainty. In M. Fishbein (Ed.), *Progress in social psychology* (pp. 49-72). Erlbaum, Hillsdale, NJ.