

Qualitative Spatial Reasoning: Framework and Frontiers

Kenneth D. Forbus
Qualitative Reasoning Group,
The Institute for the Learning Sciences Northwestern University,
1890 Maple Avenue, Evanston, IL, 60201, USA

Abstract: Spatial reasoning is a diverse topic: what might different spatial tasks have in common? One task where substantial progress has been made is qualitative spatial reasoning about motion. Unlike qualitative dynamics, purely qualitative spatial representations have not proven fruitful. Instead, a diagrammatic representation appears to be necessary. This paper begins by outlining the *Metric Diagram/ Place Vocabulary* (MD/PV) model of qualitative spatial reasoning, illustrating its power with via two example systems: FROB, a system which reasoned about motion, and CLOCK, a system which analyzed fixed-axis mechanisms. We believe this model is applicable beyond simply reasoning about motion. We suspect that (1) some form of metric diagram is a central unifying factor in all spatial reasoning tasks and (2) for human spatial reasoning, the metric diagram is part of, or at least grounded in, our perceptual apparatus. In this spirit, we identify three other kinds of spatial reasoning tasks as research frontiers where substantial progress might also be made, and pose six challenge problems to serve as milestones. The frontiers are (1) deriving system function from concrete structural descriptions (2) representing and reasoning about spatially distributed systems and (3) explicating the role of visual perception and recognition in spatial reasoning.

1. Introduction

Given the context of this book, we take the importance of spatial reasoning as given. While Cognitive Science is far from a comprehensive understanding of spatial reasoning, there has been substantial progress in specific subproblems, as the papers in this book demonstrate. The purpose of this chapter is to look at one of those areas, qualitative spatial reasoning about motion, to show the ideas which worked and to examine how they might be applied more broadly. We begin by outlining the *Metric Diagram/Place Vocabulary* (MD/PV) model of qualitative spatial reasoning (Section 2) illustrating it through two examples (FROB and CLOCK). Section 3 concludes by discussing three frontiers in spatial reasoning, highlighting how extensions of these ideas might apply, and poses some specific challenge problems to serve as milestones.




Figure 1: Qualitative spatial reasoning appears to require the derivation of task-specific qualitative representations from metric diagrams. A qualitative representation sufficient for figuring out how to assemble a bicycle must support inferences about what parts can fit together, which requires knowing relative sizes. Qualitative abstractions sufficient for this task cannot be computed for each part independently of the others, but good qualitative representations can be computed based on the set of parts. We conjecture that human spatial reasoning involves incrementally calculating such qualitative representations based on metric diagrams.

2. The MD/PV model

Diagrams and models seem inextricably linked with human spatial reasoning. Why? The wealth of concrete detail in such analog spatial representations at first might seem more than necessary for most spatial questions. Perhaps there are more abstract representations of shape and space which by themselves are sufficient for the tasks an intelligence cares about. If there are, then why don't people use them? Is it due to our highly evolved perceptual systems, which make it cheaper to measure than to infer? Have we simply not discovered them yet? Or is there some deeper reason why humans rely so much on perception for spatial tasks?

We believe there are deep reasons for human reliance on perceptual abilities for spatial reasoning, and that these reasons dominate the structure of theories for spatial reasoning in both people or machines. Our concise statement of this idea is the *Poverty Conjecture* (Forbus, Nielsen, & Faltings, 1987, 1991): *There is no problem-independent, purely qualitative representation of space or shape.* By “purely qualitative,” we mean to rule out representations whose parts contain enough detailed information to permit calculation or the operation of perceptual-like processing. (Alas, we do not yet have a more precise definition.) Examples of descriptions that are purely qualitative include describing a two-dimensional shape by a list of signs of curvatures for boundary segments, or stating that two objects are gears which can mesh. Examples of descriptions that are quantitative enough to permit perception-like processes are symbolic descriptions with numerical components, high-resolution arrays, and symbolic algebraic expressions.

By problem-independent, we mean that descriptions in the representation must be able to support, by themselves, a variety of tasks. Suppose we had a collection of physical parts, such as a Lego or Erector set, which can be used to build a variety of physical objects. A good problem-independent representation for such parts would allow us to compute a description for each part independently and then use the parts' descriptions to answer questions about their spatial interactions, without referring back to the original objects. But characterizing the size of, for example, beams *a priori* does not make sense --- an important property of beams is the relationship of their size to other elements in the construction kit. For a fixed set of parts, which can be combined only in highly constrained ways, such a qualitative description might be found. But adding even a single new kind of object (e.g., a differential gear box) could require going back to the original parts and computing a new descriptive vocabulary. Figure 1 illustrates.

The heavy reliance on metrical information is a general feature of spatial reasoning problems. Consider what a household robot needs to know to navigate. The simplest representation we might consider is a graph whose nodes are pieces of space the robot can be in and whose edges are paths from one piece of space to another. This abstract, topological description of space may suffice for getting the robot from one part of the house to another. But what if the robot is carrying (a) your backpack, (b) a sack of groceries, or (c) a bicycle? (Figure 2 illustrates.) In a typical house some edges of the original graph of places will no longer correspond to part of a reasonable path for the robot and its burden. This is an example of how purely qualitative representations can fail. Even adding a modicum of quantitative information (e.g., minimum width of each path) can greatly extend the usefulness of the representation. But for very subtle reasoning (e.g., getting a piano up a spiral staircase), a wealth of spatial detail is necessary to determine if a motion is feasible. This does not mean that qualitative representations are useless: far from it. For instance, given the maximum width of a burden we can incrementally compute a problem-specific qualitative representation of space which greatly simplifies path planning. The point is that for spatial reasoning, the interaction between qualitative information and more detailed information needs to be tightly coupled.




Figure 2: Even navigation problems are affected by the Poverty Conjecture. Figuring out what constitutes a reasonable topological abstraction of a physical space requires information about what objects are passing through the space, in addition to information about the space itself. Thus metric information is required for representations that are independent of specific tasks.

What does the poverty conjecture tell us about spatial reasoning? It suggests that spatial representations consist of two parts: a *metric diagram*, which includes quantitative information and thus provides a substrate which can support perceptual-like processing, and a *place vocabulary*, which makes explicit qualitative distinctions in shape and space relevant to the current task. The metric diagram can use floating-point numbers, algebra, or even high-precision arrays --- whatever it uses, there must be enough detail to support answering spatial queries by calculation, and it must be capable of supporting the construction of place vocabularies. Place vocabularies consist of *places*, contiguous regions of space where some important property (e.g., in contact with a gear, inside a well) is constant. Computing the place vocabulary according to the needs of the problem

ensures that the relevant distinctions are made. Defining the places in terms of elements in the metric diagram makes the diagram a good communication medium for diverse representations. Next we illustrate these ideas through two example systems.

2.1 Example: FROB

FROB (Forbus, 1980, 1983) reasoned about the motion of balls bouncing around in a two-dimensional diagram. Aside from predicting the specific motion of a given ball, FROB also produced on demand summaries of the eventual fate of a ball and estimates about whether two balls might collide. These problems are an important subset of the spatial reasoning tasks faced by any agent operating in a world of moving objects. For instance, one should be able to quickly figure out that two balls thrown into the same well might collide, while throwing the balls into different wells means they cannot collide, unless one of them escapes.




Figure 3: A FROB scenario. Can these balls collide? What if the right ball flies over the well and the left ball falls into it? What if both balls fall into the well? Even these simple qualitative stipulations can dramatically affect our judgements concerning their fate.

Figure 3 illustrates a typical FROB scenario. The metric diagram is initialized with a set of line segments corresponding to surfaces. Since balls are modeled as point masses, given the surfaces FROB computes a single place vocabulary that is used for all balls. Since

gravity acts vertically, free space is divided into places by vertical and horizontal lines emanating from corners. The vertical boundaries distinguish what a ball might hit when dropping downward, and the horizontal boundaries indicate important heights for reasoning about energy. For instance, if a sequence of bounces drops a ball's energy so that the maximum height it can reach is below the top free space segment defining a well, it has become trapped in that well.

FROB 's metric diagram consists of a set of symbolic descriptions with numerical attributes. For instance, a line segment includes a pointer to a physical interpretation (e.g., surface) as well as a numerical length and end points, each in turn specified via floating-point coordinates. The numerical information allows spatial queries to be answered by calculation rather than inference, while the physical information guides the selection and interpretation of results. For instance, FROB uses constraint-based simulation (Forbus & de Kleer, 1993), rather than traditional incremental-time simulation, to predict motion using numerical data. Constraint-based simulation essentially involves plugging in numerical values for analytic solutions of the original differential equations. This computation uses geometric reasoning in two ways. First, to establish what a ball is doing (flying through the air versus colliding with a surface), we must know where it is. Second, geometric reasoning is required to solve for boundary conditions. Consider a ball flying through the air with a known initial position and velocity. This information suffices to calculate a parabola (or line segment) corresponding to the dynamical constraints on its trajectory, which is then intersected with the surfaces in the scene to figure out what the ball would hit. The new state of the ball after the collision can then be computed, a new trajectory drawn, and so forth.

FROB 's place vocabulary allows assimilation of qualitative constraints and supports global reasoning about possible motions. The elements of the place vocabulary and the relationships between them provide the qualitative terms required for qualitative simulation laws for envisioning the possible motions of balls. Certain “obvious” questions about collisions can be answered by comparing envisionments: If two balls are never in the same place, for instance, they can never collide. Furthermore, an envisionment can be pruned to represent the consequences of qualitative constraints on motion (e.g., assuming that a ball doesn't go into a well) or of numerical constraints expressed geometrically (e.g., ruling out motions occurring in places above the maximum height a ball could reach given its energy).

FROB provides a good illustration of the utility of the MD/PV model of spatial reasoning. The metric diagram provides the precision required to calculate boundary conditions for detailed predictions, and provides a substrate for automatic computation of the place vocabulary. The place vocabulary supports sophisticated qualitative reasoning about motion. But even if the graph of places was provided *a priori*, many interesting questions cannot be answered without grounding the place vocabulary in the metric diagram (e.g., whether a ball is trapped in a well or ruling out collisions between balls). Thus even in FROB 's simple domain, both the metric diagram and place vocabulary are required for sophisticated spatial reasoning.

2.2 Example: CLOCK

A more complex class of motion problems concerns fixed-axis mechanisms, such as mechanical clocks. The creation of reliable mechanical clocks was a milestone in mechanical engineering; thus the development of systems which can reason about such mechanisms serves as a good milestone for qualitative spatial reasoning. One milestone, the qualitative simulation of a mechanical clock from first principles, was achieved by our group in February 1988 by the CLOCK system (Forbus, Nielsen, & Faltings, 1987, 1991) built by Paul Nielsen and Boi Faltings as part of their Ph.D. theses.

CLOCK worked on fixed-axis mechanisms which could be decomposed into two-dimensional interactions. Given a CAD-like description of the parts of a mechanism and their degrees of freedom, CLOCK computed a place vocabulary based on configuration space. Configuration space is the appropriate basis for the place vocabulary because connectivity is central to kinematic state. Faltings developed an elegant characterization of such places and algorithms for computing them (Faltings, 1987, 1990). Instead of developing a single, massive high-dimensionality vocabulary, his algorithms created place vocabularies for pairs of parts that could interact, and used elements from these vocabularies to define kinematic states for the whole mechanism. The metric diagram plays a key role in this composition process, since defining consistent combinations of places sometimes requires projecting constraints from one vocabulary to another (and thus introducing new distinct places) when two vocabularies share a part.

CLOCK's metric diagram also played a key role in keeping envisioning tractable. Any diagram has noise, and the well-known sensitivity of kinematics to the details of surface shape means that smoothing the contours of a part in isolation is inappropriate. Nielsen realized that filtering at the level of the place vocabulary allows the interaction of the parts to be taken into account, and developed algorithms for removing "small" places and merging places that were "very close." The resulting simplification of the place vocabulary was dramatic: The number of potential kinematic states dropped from over 10,000 to 58 (Nielsen, 1989).

For qualitative reasoning about the dynamics of motion, Nielsen developed a qualitative vector algebra, and showed how shapes should be decomposed qualitatively to define robust notions of mechanical constraint (Nielsen, 1990). The requirements of reasoning about mechanical constraint again illustrate the need for interactions between qualitative and quantitative representations: The appropriate qualitative description of an object for figuring out how it can move if pushed depends not just on its shape, but also the qualitative direction to its center of rotation.¹

¹ A simple demonstration: The center of rotation for a book on a table is its center of mass. Pick a side of the book and try pushing it on the corners versus the middle. Notice that the sign of its rotation differs in each case, providing a criteria for decomposing its shape qualitatively. Now drive a nail through the book into the table, somewhere near a corner of the book. The center of rotation will now be different, and hence a different qualitative decomposition of the book's shape is required.

2.3 Supporting Evidence

The ideas described above are not the only work which supports the claim that the combination of metric diagrams and place vocabularies is necessary for spatial reasoning about motion (c.f. Joskowicz, 1987, Joskowicz & Sacks, 1991; Narayanan & Chandrasekaran, 1991). A more complete examination of the relevant literature and its relation to this model can be found in (Forbus, Nielsen, & Faltings, 1991). As the household robot example suggests, we believe the MD/PV model is relevant to path-planning problems. Furthermore, the widely reported use of imagery in scientific and engineering reasoning (c.f. Tweney, 1990) suggests that similar constraints operate in other forms of spatial reasoning as well. While powerful, self-sufficient, purely qualitative representations for shape and space would in many ways be desirable, none has yet been found and there are good reasons to suspect that they simply do not exist. If this is correct, the necessity of using some form of metric diagram may be the central unifying factor in spatial reasoning.

Many interesting physical systems, including internal combustion engines and pumps, can only be understood by comprehending the interactions of fluids and solid objects through motion. Recently Hyeonkyeong Kim (Kim, 1993) has developed a set of qualitative representations and reasoning techniques that are sufficient for analyzing such examples. She developed a qualitative account of linkages that can envision up to four-bar linkages, showing that representing relative lengths of links, when combined with a qualitative representation of angle that includes both what quadrant it is in and relative inclination with respect to other links, suffices to solve this problem (Kim, 1992a). She developed a qualitative method of computing streamline diagrams for laminar fluid flow, capable of generating explanations about how the Bernoulli effect generates lift over an airplane wing (Kim, 1990). Complex processes such as the combustion in a cylinder that drives a piston downward are typically approximated as instantaneous changes, so she developed a qualitative theory of impulses, showing how their effect could be modeled via structural changes in envisionments (Kim, 1992b). Finally, figuring out how liquids can move inside containers that contain pistons and valves requires describing fluids in terms of their potential surface contacts. Kim developed such a representation for fluids, exploiting Nielsen's work on representing shapes and surfaces (Kim, 1992c). Kim's research is a large step towards a comprehensive qualitative mechanics.




Figure 4: An example analyzed by SKETCHY. Comparative analysis can be carried out by visual manipulations on diagrams. Here, the effect of superheating in a steam plant is represented graphically by a change in the position of point 3 in this temperature/entropy diagram. SKETCHY's visual routines enable it to compute that this change increases the area under the curve, and its interpretation of these geometric changes in physical terms allows it to conclude that the efficiency of the system is thus increased.

What about the connection of spatial reasoning to perception? As vision researchers begin to grapple with higher-level perception, more attention has been placed on the role of task constraints in visual processing. For example, Ullman has postulated a *visual routines processor* which is used for recognition tasks and is programmable by higher systems (Ullman, 1985.) An interesting conjecture is that the human metric diagram is in fact the same thing as the visual routine processor. Approaching these issues from opposite sides of the cognition/perception borderland has meant that each effort focuses on different tasks, but this conjectured identity raises some interesting avenues for joint empirical inquiry (c.f. Chapman, 1991; McDougal & Hammond, 1992).

One task for which visual routines are likely to be used is in interpreting line graphs, such as the temperature-entropy diagrams commonly used in thermodynamics to reason about the properties of power plants and refrigerators. Yusuf Pisan (Pisan, 1994) has shown how visual processing on metric diagrams can be used to reason with such pictures. His system, SKETCHY, takes as input a diagram, expressed as line segments (curves are approximated by many short line segments). SKETCHY uses the labels associated with the diagram elements to infer information about physical properties. For example, SKETCHY uses vision algorithms to estimate curvature, and thus can extract qualitative proportionalities from graphs. An especially interesting and novel aspect of SKETCHY is

that it demonstrates Pisan's observation that an important use of diagrams in human reasoning is *comparative analysis* (Weld, 1990). Weld's work on comparative analysis argued persuasively that it is an important style of human reasoning about physical systems, but also showed that there were strict limits on how well it could be supported by purely qualitative representations. Pisan argues that the availability of metric information in diagrams makes them a natural representation for comparative analysis. For example, in Figure 4 the physical changes caused by superheating the working fluid in a Rankine cycle power plant are depicted in a diagram (from Whalley, 1991). SKETCHY uses its visual routines to calculate that this change increases the area under the curve, which it then interprets physically as an increase in the efficiency of the system. To date SKETCHY has analyzed over 60 diagrams from thermodynamics and economics textbooks, supporting these and other inferences.

3. Frontiers and Challenge Problems

Progress continues in qualitative spatial reasoning about motion, based on the MD/PV model, and this aspect of the problem of spatial reasoning is by no means completely tamed. However, there are many other aspects of spatial reasoning that also deserve attention, and seem ripe for substantial progress as well. Furthermore, the recent improvements in floating point and graphics hardware, along with impressive increases in memory speeds and capacities, means that technological barriers for computational experiments are falling: Spatial reasoning may be one of the few areas in AI where our progress has become more limited by ideas than by technology.

In this spirit, the remainder of this paper proposes several challenge problems as possible focal points for research on spatial reasoning. Our end goal as Cognitive Scientists remains unchanged, of course: a comprehensive computational account of spatial reasoning, with both an empirically tested explanation for human spatial reasoning capabilities and practical representations and algorithms for mechanical reasoners. The problems proposed here are meant to be milestones, in that their solutions would count as substantial progress even though they appear (at least from our current perspective) to be somewhat simpler. In each case we first sketch the general task and why it is important, and then suggest some challenge problems.




Figure 5: A spring-loaded reducing valve. This picture, from a U.S. Navy training manual, is the original source of the abstract reducing valve example widely used as an example in qualitative physics. Many interesting aspects of physical reasoning are involved in translating from this picture into the abstract representations that most qualitative physics work starts with.

3.1 Deriving function from real structure

An important engineering skill is interpreting blueprints and schematics, deriving from the description of the structure of the system its intended function. Previous AI work on this problem has focused on analog electronics (de Kleer, 1984), an unfortunate choice from our perspective since that domain is designed to allow geometric considerations to be mostly ignored. The structural descriptions used in more recent qualitative physics research mostly remains quite abstract, ranging from specifying a system in terms of qualitative equations to describing idealized entities such as containers and abstract fluid paths (Kim's work, as outlined above, is an exception). Such descriptions may be good

candidates for the output of systems which assimilate information from more standard sources, but the problem of deriving abstract structural descriptions from more primitive descriptions itself deserves study. For example, a standard example in qualitative physics is a kind of spring-loaded reducing valve. Figure 5 shows the original spring-loaded reducing valve which motivated subsequent work. The process of creating a usable mental model from this geometric description and some supporting text does not seem easy.

Two major issues in such reasoning seem to be (a) how to move from primitive descriptions of shape, material composition, and connectivity to a more abstract interpretation of the parts of the device suitable for other reasoning tasks and (b) how to manage the complexity of the structural descriptions so created (c.f. Liu & Farley, 1991). These issues suggest two challenge problems:

Challenge Problem 1: Develop a system which can answer questions about the possible behaviors and uses of pneumatic control components, given the kind of drawings and annotations provided in a typical textbook.

Challenge Problem 2: Develop a system which can, given the blueprints of a complex system such as an automobile engine, derive a decomposition of it into subsystems and provide abstract functional models of the subsystems and overall operation of the system.

3.2 Reasoning about spatially distributed systems

Qualitative dynamics has focused on lumped-parameter systems. Many phenomena, including heat transfer, aerodynamics, and fluid dynamics are best modeled as spatially distributed systems. Developing the qualitative version of partial differential equations requires deeper and more flexible integration of spatial reasoning with dynamics. The potential applications are very important: A key problem in today's use of finite element analysis in problems such as computational fluid dynamics is quantizing shapes and space to provide accurate results with the least computational effort. Programs which could automatically produce good meshes would be quite valuable.

One key issue seems to be finding appropriate constraints on place vocabularies. What are the analogs to CLOCK's decomposition of shapes along centers of rotation, or FROB's decomposition of free space along vertical and horizontal axes? The place vocabularies must both support computations with more detailed data and support explanations of overall qualitative features of behavior.

Challenge Problem 3: Develop a system which can automatically set up meshes for finite element analysis programs, given a CAD description of the physical system to be analyzed.²

² Since this paper originally appeared in the AAAI Spring Symposium, some progress has been made in this area (c.f. Yao & Gelsey, 1993).

Challenge Problem 4: Develop a system which can, given a sequence of weather maps for a region, provide a consistent qualitative explanation of the atmospheric behavior during that period, or detect if the sequence of maps is inconsistent.

3.3 Perception in spatial reasoning and learning

In many applications of spatial reasoning the inputs are self-evident (e.g., CAD files or blueprints or weather maps). For cognitive simulation the problem of what the inputs are is a subtle methodological question. Inappropriate assumptions about input can trivialize important problems or divert us into solving non-problems. The psychological aspect of the MD/PV model is the idea that the Metric Diagram is in fact part of our perceptual system, programmed in part according to task and in part bottom-up. Consequently, questions about the nature of high-level perception and the relationship between perception and cognition are crucial to cognitive simulations of spatial reasoning. For instance, in cognitive developmental psychology there is a fascinating set of theories and empirical findings on children's acquisition of ideas concerning motion, objectness, and causality (c.f. Baillargeon & DeVos, 1991, Spelke, 1991). Furthermore, learning and teaching often involve interaction through mixtures of text, speech, and diagrams (Shrager, 1990).

Challenge Problem 5: Simulate the acquisition of commonsense notions of objects, space, or causality, including an artificial vision system as part of the input processing.

Challenge Problem 6: Develop a system which can be taught a qualitative theory of a phenomena, such as aerodynamics or the greenhouse effect, via an interactive, mixed-medium dialog.

4. Acknowledgements

The research on reported here was supported by the Computer Science Division of the Office of Naval Research. I thank Dedre Gentner, Ron Ferguson, Yusuf Pisan, and David Foster for useful discussions and comments.

5. Bibliography

Baillargeon, R. and DeVos, J. "Object permanence in young infants: Further evidence." *Child Development*, **62**, 1227--1246, 1991.
Chapman, D. *Vision, Instruction, and Action*, MIT Press, 1991.
de Kleer, J. "How Circuits Work", *Artificial Intelligence*, **24**, 1984.

Faltings, B. "A theory of qualitative kinematics in mechanisms", University of Illinois at Urbana-Champaign, Department of Computer Science Technical Report No., UIUCDCS-R-86-1274, May, 1986

Faltings, B. "Qualitative kinematics in mechanisms", *Artificial Intelligence* **41**(1), June, 1990

Forbus, K. "Spatial and qualitative aspects of reasoning about motion", Proceedings of the National Conference on Artificial Intelligence, Palo Alto, California, August, 1980

Forbus, K. "Qualitative reasoning about space and motion" in Gentner, D. and Stevens, A. (Eds.), *Mental Models*, LEA associates, inc., 1983.

Forbus, K. and de Kleer, J. *Building Problem Solvers*, MIT Press, 1993.

Forbus, K., Nielsen, P., and Faltings, B. "Qualitative kinematics: A framework" Proceedings of IJCAI-87, Milan, Italy, August, 1987.

Forbus, K., Nielsen, P. and Faltings, B. "Qualitative Spatial Reasoning: The CLOCK Project", *Artificial Intelligence*, **51**(1-3), October, 1991.

Joskowicz, L. "Shape and function in mechanical devices", Proceedings of the Sixth National Conference on Artificial Intelligence, August, 1987

Joskowicz, L. and Sacks, E. "Computational kinematics", *Artificial Intelligence*, **51**(1-3), October, 1991.

Kim, H. "Qualitative reasoning about the geometry of fluid flow", *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, Cambridge, MA, July, 1990.

Kim, H. "Qualitative Kinematics of Linkages." In Faltings, B. and Struss, P. (Eds.) *Recent Advances in Qualitative Physics*, MIT Press, 1992.

Kim, H. "Augmenting qualitative simulation with global filtering", *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, 1992.

Kim, H. "Extending the contained-stuff ontology with geometry." *Proceedings of the Sixth International Qualitative Physics Workshop*, 1992.

Kim, H. "Qualitative reasoning about fluids and mechanics", Technical Report #47, The Institute for the Learning Sciences, Northwestern University, December, 1993.

Liu, Z. and Farley, A. "Tasks, Models, Perspectives, Dimensions", Proceedings of the Fifth International Workshop on Qualitative Reasoning about Physical Systems, May, 1991, Austin, Texas.

McDougall, T. and Hammond, K. "A recognition model of geometry theorem proving", *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, 1992.

Narayanan, N.H., and Chandrasekaran, B. "Reasoning visually about spatial interactions", *Proceedings of IJCAI-91*, Sydney, Australia, 1991.

Nielsen, P. "A qualitative approach to mechanical constraint", Proceedings of the National Conference on Artificial Intelligence, August 1988; also in *Qualitative Reasoning about Physical Systems*, Daniel S. Weld and Johan deKleer Eds., Morgan Kaufman Publishers, Inc., 1990

Nielsen, P. "The role of abstraction in place vocabularies", Proceedings of the Conference of the Cognitive Science Society, 1989.

Pisan, Y. "Visual reasoning about physical properties via graphs", to appear in *Proceedings of QR94*, Nara, Japan, June, 1994.

Shrager, J. "Commonsense Perception and the Psychology of Theory Formation", in Shrager, J. and Langley, P. *Computational Models of Scientific Discovery and Theory Formation*, Morgan-Kaufmann, 1990.

Spelke, E.S. "Physical knowledge in infancy: Reflections on Piaget's theory", in S. Carey and R. Gellman (Eds.) *The Epigenesis of Mind*, New Jersey, LEA, Inc., 1991

Steinberg et.al.

Tweney, R. D. "Five questions for computationalists", in Shrager, J. and Langley, P. *Computational Models of Scientific Discovery and Theory Formation*, Morgan-Kaufmann, 1990.

Ullman, S. "Visual Routines", in Pinker, S. (Ed.) *Visual Cognition*, MIT Press, Cambridge, MA, 1985.

Weld, D. *Theories of Comparative Analysis*, MIT Press, 1990.

Whalley, P. *Engineering Thermodynamics*, Oxford University Press, 1991.

Yao, K.T. & Gelsey, A. "Spatial reasoning for intelligent control of numerical simulators", *Proceedings of QR93*, May, 1993.