

Qualitative and Quantitative Reasoning over Physics Textbook Diagrams

Maria D. Chang, Jon Wetzel, Kenneth D. Forbus
Qualitative Reasoning Group
Northwestern University
Evanston, IL 60208

Abstract

Diagrams are frequently used in problem solving to convey spatial and conceptual information. They can guide problem solving in domains that require both qualitative and quantitative reasoning about causal systems. One such domain is physics, which requires qualitative and quantitative reasoning for conceptual mastery. However, almost no physics problem solving systems use diagrams or hand drawn sketches as input. We present a system that uses qualitative process theory and qualitative mechanics to solve diagrammatic ranking exercises from a physics textbook. The combination of quantitative and qualitative reasoning over sketches enables the software to understand physical situations in human-like ways that are similar to humans. The application of these reasoning techniques to sketched physical systems may be useful for educational software in physics and engineering design.

1 Introduction

Humans are constantly engaging in qualitative reasoning about physical situations. We can interact with, and make predictions about, the physical world in the face of limited information and uncertainty. Physics educators have argued that qualitative, concept-based instruction helps students select problem solving strategies [Leonard *et al.*, 1996]. This has led many educators to make qualitative physics the focus of early physics education, under the premise that qualitative, conceptual understanding precedes formal proficiency. Thus, educational software that is designed to understand physical concepts in ways similar to humans must have the ability to reason about qualitative, incomplete descriptions of physical systems.

Representations of physical systems are also enriched by sketches and diagrams. Externalized diagrammatic representations have the advantages of reducing working memory load and allowing spatial inferences to be computed with greater ease [Larkin and Simon 1987]. The process of sketching is also a hallmark of classroom activities in spatial domains. Sketching is a critical step in the design of physical systems and sketching behavior may be indicative of

expertise (e.g. biology [Kindfield, 1992] and geological processes [Jee *et al.*, 2009]). Arguably, any educational software that deals with physics understanding and lacks sketches or diagrammatic representations is incomplete.

One approach for solving everyday physics problems is via case-based reasoning. In [Klenk *et al.*, 2005], analogical retrieval was used to recall and apply causal models from past experiences to novel problems. Causal information from previous situations along with qualitative mechanics inferences were used to solve comparative analysis problems from the Bennett Mechanical Comprehension test. This involved comparing two quantities across two or three different, but structurally similar, scenarios.

The conceptual physics ranking exercises described in this paper are very similar to the comparative analysis questions examined by [Klenk *et al.*, 2005]. Our approach, however, uses first principles analysis to form causal models for each new problem, rather than using analogy to recall past causal models. Our system is also designed to resolve quantity differences between any number of different, but structurally similar, scenarios.

In this paper, we describe how qualitative spatial reasoning, qualitative mechanics and qualitative process theory can be used to solve diagrammatic ranking exercises from a physics textbook. We begin by briefly reviewing CogSketch, qualitative mechanics, and QP theory. Then we describe our approach, and illustrate its utility by showing its performance of a set of ranking problems. We close with a discussion of related and future work.

2 CogSketch

CogSketch is an open-domain sketch understanding system [Forbus *et al.*, 2011]. It collects and understands user drawn sketches by modeling the perceptual and spatial understanding that humans use when sketching. CogSketch uses qualitative spatial representations, embracing the informal (and robust) nature of human to human sketching. Examples of qualitative spatial relations include positional relations (e.g. above, rightOf) as well as topological relations (i.e. region connection calculus [Cohn *et al.*, 1997]).

The spatial information in CogSketch is tied to conceptual information via *conceptual labeling*, which allows users to label their drawn elements with concepts from the Open-

Cyc knowledge base (KB)¹. Quantities can be denoted using *sketch annotations*, which associate specific quantities (e.g. height, gravitational force, etc.) with drawn objects. Thus, CogSketch understands that the sketched objects represent entities and quantities with corresponding properties. This means that CogSketch does not rely on sketch recognition. This is a deliberate design decision based on the observation that humans rarely sketch things neatly enough to be reliably recognized; instead, they typically rely on natural language or gesture to indicate what their drawing represents. Additionally, CogSketch is designed for open-domain sketch understanding. Today's sketch recognition algorithms are limited to small to medium sized domains that are completely specified in advance.

The spatial and conceptual information gathered by CogSketch can be reasoned about using the structure mapping engine (SME) [Falkenhainer *et al.*, 1989], which is based on the structure-mapping theory of analogy [Gentner, 1983]. SME can be used to compare spatial representations to highlight qualitative similarities and differences.

CogSketch has been used to model spatial problem solving (e.g. geometric analogies and Raven's Progressive Matrices) [Lovett & Forbus, 2010] and to collect sketching information for psychological experiments [Jee *et al.*, 2009]. In education, CogSketch has been used as a platform for software-based sketch worksheets [Yin *et al.*, 2010] and as a tool to help engineering students learn to communicate about their design sketches [Wetzel & Forbus, 2010].

3 Qualitative Mechanics

As described in Wetzel and Forbus [2008], we use a model of qualitative mechanics (QM) that is based on the work of Nielsen [1988] and Kim [1993]. Like their work, our model of qualitative mechanics works in 2-dimensional space and is able to represent forces acting on/between objects, calculate the direction of net force and motion of an object and predict the behavior of future states. Unlike their work, our system can take hand drawn sketches as input. Our qualitative mechanics reasoning facilities are built in to CogSketch.

Qualitative mechanics allows CogSketch to understand sketched mechanical systems and has been used in educational software intended for engineering design education [Wetzel & Forbus, 2010]. However, the previous implementations did not deal with ordinal relationships between quantities. This is one of the reasons why qualitative process theory (reviewed next) was used in our approach for solving conceptual physics ranking exercises. In this work we extend qualitative mechanics by taking qualitative velocity and net force vectors (which QM could calculate before) and making statements about the magnitude of those vectors. Thus, the same representations which worked in QM can be used in this new system.

4 Qualitative Process Theory

Qualitative process (QP) theory is a representational system that organizes physical phenomena around *physical processes*, which impose causal relationships on *continuous quantities* [Forbus 1984]. Physical processes are the sole mechanism of change in QP theory. Examples of physical processes include things like liquid flow, heat flow and boiling. Every process can have logical *consequences* and *direct influences* on continuous quantities (e.g. amount of liquid, amount of heat, amount of steam). A causal relationship can be an *indirect influence* as well. Indirect influences can be used to describe functional dependence between continuous quantities. To say that one quantity increases monotonically with another, all else being equal, we say that the first is *qualitatively proportional* (q_{prop}) to the second. To say that one quantity decreases monotonically with another, all else being equal, we use a negative qualitative proportionality (q_{prop-}) between the first quantity and the second. Qualitative proportionalities can be combined to form more complete causal models for quantities. For example, the functional dependence between acceleration and force and mass as described in Newton's second law, $F = m * a$, can be summarized as:

$$\begin{aligned} & (q_{prop} \ a \ F) \\ & (q_{prop-} \ a \ m) \end{aligned}$$

Here we use only two of the basic inferences of QP theory, namely model formulation and influence resolution². This requires a scenario description and one or more domain theories. The scenario description is a predicate calculus representation of the objects and relationships that we want to reason about. The domain theory defines the processes, quantities and influences that can be used to formulate causal models about the scenario and make inferences. Model formulation [Falkenhainer & Forbus, 1991] is the process of analyzing the scenario description and determining what qualitative models from the domain theory are applicable. If a model is applicable to the scenario and its preconditions are met, then its consequences are inferred to the scenario model. These inferences often include causal relationships between quantities. Once a qualitative causal model is formed about the scenario, influences on quantities can be resolved. This allows knowledge about quantities to be propagated along a causal chain. For example, if we know that acceleration increases with force and we know that force is increasing, then we know that acceleration is increasing as well.

These building blocks are powerful enough to create qualitative models and simulations of a wide range of physical and conceptual phenomena. Qualitative models using QP theory can represent physical phenomena like fluid dynamics and thermodynamics as well as conceptual phenomena like describing how a credit card works. The basic building blocks remain the same across domains, making

² Limit analysis is not needed because, for the ranking problems in [Hewitt 2010], each situation being compared consists of only a single qualitative state.

¹ www.openencyc.org

QP theory a powerful representation system for teaching complex models in a simple manner.

5 Approach

Our approach for solving conceptual physics ranking exercises can be broken down into four major steps. First, the problem scenario is sketched by hand into CogSketch, which captures spatial, conceptual and quantity information about the sketch. Next, QM is used to determine net forces on objects, surface contacts and possible future states. Then, QP theory is used to detect causal models and continuous processes in the problem scenario. It is here that qualitative causal relationships enable reasoning about inexact quantities. Lastly, differential qualitative analysis results in ordinal relationships between quantities in different scenarios. This last step is what allows the system to rank the scenarios along a particular quantity.

5.1 Conceptual Physics Ranking Exercises

To evaluate the ability to solve conceptual physics problems, we attempted to solve ranking exercises from Hewitt's [2010] *Conceptual Physics* textbook. The textbook is divided by topic into eight parts. We focused on the first part of the book: mechanics. In part one, there are four chapters with 27 ranking exercises in total. These exercises tap knowledge of forces, acceleration, velocity, friction, tension and kinetic and potential energy. As a starting point, we selected the ranking exercises that covered net force, net velocity and tension, which make up 12 of the 27 exercises. The results reported below describe the performance of our system on those 12 exercises.

Ranking exercises are physics problems in which two to four similar physical scenarios are presented and the student is asked to rank the value of a quantity across them. Figure 1 shows an example from [Hewitt 2010] in which two people are standing on a scaffold supported by two ropes. The goal is to rank the tension in the left rope in the three situations from greatest to least.

Many of the ranking exercises are like Figure 1, highly




Figure 1. Ranking problem example: rank the tensions in the left rope from greatest to least. The answer is $C > B > A$.




Figure 2. Ranking problem example: rank the net-force on the boxes. The answer is $D > A = B = C$.

qualitative in nature, which makes QP theory well suited to the task. Some of the problems however are more quantitative in nature, like Figure 2. For problems like this, where a qualitative causal model is not necessary, we determine the quantities at each step using an extended version of qualitative mechanics and then compare the values directly to reach the answer. The extensions to QM include a representation for one object moving in a large moving object (e.g. a person on a train, a boat in a river) and finding net vector magnitudes (previously it was only direction).

5.2 Problem Scenario Descriptions

Textbook diagrams describing ranking problems contain multiple parts, each of which must be represented and reasoned about separately before combining the results to solve the problem. CogSketch handles such situations naturally, since its sketches are further decomposed into *subsketches*. The “whole thing” is represented on the *metalayer*, where each subsketch is treated as an object. Each subsketch has its own reasoning context, i.e. a Cyc microtheory. The microtheory for a subsketch contains facts about that subsketch including spatial relationships automatically computed by CogSketch, conceptual labels chosen by the user, and values for quantities entered by the user (e.g. the forces and masses in Figure 2).

Consequently, sketching a ranking problem consists of making one subsketch per portion of the problem. The entire sketch itself constitutes the problem.

5.3 QP domain theory

We created a domain theory for describing two dimensional mechanics. This domain theory consists of axioms, rules, and qualitative *model fragments* which define the processes, quantities, and influences used in QP theory. Quantities we modeled included net force, position, velocity, acceleration, inertia, gravitation, normal forces, friction, mass, tension and distance between objects.

The continuous processes modeled include motion (with or without friction) and acceleration. These processes have direct influences on position and velocity respectively. Physical model fragments encode particular phenomena that can occur in situations (e.g. a scaffold hanging by two ropes), an object with inertia and an object in free fall.

Figure 3 illustrates a model fragment that represents the relationships between tensions that arise from two ropes supporting something. The model fragment has four participants: two hangers (ropes), one hanging thing (a solid thing) and a set of all things supported by the system. In order for this model fragment to be instantiated for a given physical scenario, certain constraints must be met. In this case, the hanging thing must hang from both hangers. Additionally, the set of all things supported by the system includes the hanging thing and all the things that the hanging thing supports. Using the problem in Figure 1 as an example, this set would include the scaffold and the two persons that are supported by the scaffold. The vertical tensions depend on the downward forces of all the hanging things in the system. The actual values of those forces are unknown,

but it can still be inferred that the vertical tensions are dependent on the location of the center of mass of all the hanging things. More specifically, the closer a rope is to the center of mass of all the hanging things, the greater that rope's vertical tension. This functional dependence is captured by the negative qualitative proportionalities ($qprop^-$) in the model fragment.

```

Model fragment HangingSomething-ByTwo
Participants:
  firstHangerOf: ?left, a Rope
  secondHangerOf: ?right, a Rope
  hangingThingOf: ?scaffold, a Solid
    (hangsFrom ?scaffold ?left)
    (hangsFrom ?scaffold ?right)
  hangingThingsOf: ?allHangingThings, a Set
    (evaluate ?allHangingThings
      (SetUnionFn ?scaffold
        (TheClosedRetrievalSetOf ?thing
          (supportedBy ?thing ?scaffold))))
Conditions:
  (hasQuantity ?right (YTensionFn ?right))
  (hasQuantity ?left (YTensionFn ?left))
Consequences:
  (qEqualTo (XTensionFn ?left)
    (XTensionFn ?right))
  ( $qprop^-$  (YTensionFn ?left)
    (DistanceFn ?left ?allHangingThings))
  ( $qprop^-$  (YTensionFn ?right)
    (DistanceFn ?right ?allHangingThings))

```

Figure 3: Example model fragment describing a scaffold hanging by two ropes. The vertical tension in each rope is negatively qualitatively proportional to the distance between the ropes and the group of things it supports.

As usual, multiple model fragments can be active at once and that is typically necessary to make meaningful inferences about a scenario.

5.5 Differential Qualitative Analysis

To solve a ranking exercise we must compare the value of a quantity in one subsketch to that in another and figure out the ordinal relationship between the two. This process is called differential qualitative analysis (DQA) [Weld, 1988]. The input to our DQA algorithm consists of a sketch, the name of the quantity being ranked, and the object to which that quantity belongs. First, QP analysis is performed on each subsketch in the sketch. This begins with model formulation and, if model formulation is successful, continues with influence resolution. When the QP analysis is successful, it means that one or more model fragments were applicable and active in the scenario. When a model fragment is active, its consequences are inferred. Since these consequences can be ordinal or functional relationships, they influence other quantities in the scenario. Influence resolution determines the causal chain between quantities, which is used for propagating DQ values through the quantities of a situation.

When the QP analysis is unsuccessful, we assume that annotations and qualitative mechanics suffice (e.g. Figure 2) and proceed to try to find the quantity values directly. Concrete quantities can be inferred from a subsketch in two ways: via sketch annotations or via quantitative analysis of the sketching space. An example of using a sketch annotation to derive a quantity is illustrated in Figure 2, where each force arrow is a sketch annotation with a particular quantity associated with it. In this case, the quantity can be looked up in the sketch's knowledge (i.e. microtheories). This is the easiest method for finding quantities, but sketch annotations with specific quantities can only be used if the problem provides us with that information. For problems that are more qualitative in nature, like the hanging scaffold problem in Figure 1, we have to infer some visual quantities by using the x and y coordinates of the sketching space. This is how distance is computed between objects in the hanging scaffold problem. Although the actual values are not needed for calculations, the ordinal relationships between them are needed to solve the problem.

To represent ordinal relationships between quantities in different subsketches, we must first compare the subsketches to each other to determine which objects and quantities are in correspondence. SME is used for this comparison because it puts objects into correspondence based on common relational structure. For example, in Figure 1, when we refer to the tension in the left rope, we know that we must compare the tension in the left rope in scenario A, to the tension in the corresponding rope in scenarios B and C.

The differential qualitative value relationship itself is represented using a relationship called $dqValue$:

```
(dqValue ?quantity ?mapping ?value)
```

Where $?quantity$ is a named quantity like tension, $?mapping$ is an analogical mapping between two scenarios, which puts objects into correspondence, and $?value$ is one of four values: IncreasedDQ, DecreasedDQ, UnchangedDQ or AmbiguousDQ.

The $dqValue$ of a particular quantity between two subsketches can be derived in one of two ways:

- directly – The $dqValue$ can be derived directly if the quantity is readily available in the sketch as an assumption, a sketch annotation or if the quantity can be inferred using QM or a quantitative analysis of the sketching space.
- via influences – The $dqValue$ can be derived via influences if the quantity cannot be derived directly but it has causal antecedents that are known.

The challenge with DQA, then, is formalizing the mechanisms for deriving quantities based on their causal antecedents. To do this, the dependency order that is created by influence resolution in QP analysis is used to determine what quantities should be solved for based on a particular goal quantity.

Lastly, the results of the DQA (increased, decreased, unchanged or ambiguous) are used to determine the ordering of the subsketches, which is passed as the solution to the exercise.

8 Results

The system correctly ranked 9 out of 12 net force, net velocity and tension questions in part one (mechanics) of the textbook.

<i>Concept</i>	<i>Number of Test Cases</i>	<i>Number Solved</i>
<i>Net Force</i>	4	3
<i>Net Velocity</i>	4	3
<i>Tension</i>	4	3

These results are based on first principles approach to problem solving and suggest that using QP theory and QM to derive qualitative differences between scenarios is feasible and promising.

The problems that were not solved (one from each concept category) suggest areas for improvement. The net force problem that failed was a question about weight on different planets. It required the system to infer gravitational force of objects based on their mass alone. In other words, it required the assumption of gravitational force without an explicit force arrow. In other net force problems (e.g. Figure 2) all the forces needed to solve the problem were represented explicitly with force arrows. However, the system cannot assume gravitational forces in all cases because this would cause other problems to be reasoned about incorrectly. Instead, we will need to extend the system so that it can determine when gravity should be assumed and when it shouldn't, perhaps by using the textual problem statement as a guide. The net velocity problem that failed involved inferring the amount of time (which relied on the vector sum of velocities) it took an object to reach a destination. In this case, it was not the qualitative vector summation that failed, but the link from qualitative vector to time. Lastly, the tension problem that failed involved inferring the ordinal relationship between the net tension of two ropes based on the ordinal relationships between the x and y components of tension. At the time of this experiment, this type of inference was not supported. In each of these cases, it appears that extending the domain theories to make inferences about more quantities (e.g. time) would be enough to arrive at the correct solution. An important feature to consider is the extent to which these inferences are made automatically. In order for this system to scale, it is important for additional inferences to be computed on an on-demand basis. We are currently extending our domain theory to overcome these limitations.

9 Related Work

Diagrams are essential in scientific domains such as physics and biology: For example, [Chaudhri *et al.*, 2009] found that 48% of physics problems on AP exams had diagrams, and in 58% of those (28% of the total) the problem could not be solved without extracting information from the diagram. The difficulty of reasoning with diagrams can be seen from the fact that despite this analysis, their AURA system for question-answering in scientific domains does

not currently include any diagrammatic reasoning capabilities. Similarly, while the Atlas-Andes [Rose *et al.*, 2001] and Autotutor [Graesser *et al.*, 2003] conceptual physics tutors can include diagrams to be presented to the student as part of their problem specification, the diagrams are not interpreted by the software itself.

Several efforts have been made to build systems that can understand diagrams in the context of physics problems. BEATRIX [Novak & Bulko 1990] combines problem information from text and diagrams. Its diagrams were created by a specialized drafting program, and hence were noise-free drawings rather than hand-drawn sketches. BEATRIX performed only minimal spatial reasoning, focusing on purely quantitative physics problems. Similarly, the Figure Understander [Rajagopalan, 1995] was limited to precisely understood diagrams, and did only qualitative reasoning about possible motions. By contrast, Lockwood's MMKCAP [Lockwood & Forbus, 2009], combined natural language understanding and information from hand-drawn sketches to read a chapter on levers and answer questions at the end of the chapter. Unlike this experiment, those questions did not involve comparing quantities across different scenarios. Lastly, [Klenk *et al.*, 2005] is very similar to this work in its use of differential qualitative analysis on sketches, but our qualitative mechanics is more advanced, the ranking problems are slightly harder than the straight comparison problems explored there, and our approach is to use first-principles reasoning, rather than case-based reasoning.

10 Conclusion

We have demonstrated that it is possible to use QM and QP theory over sketched physics diagrams to solve conceptual physics ranking exercises. Educational software that helps people improve their conceptual understanding of physics will need human-like qualitative reasoning abilities, in order to handle the range of problems that they must help students with. The progress reported here represents a step in that direction.

Much remains to be done, of course. Next, we plan to extend the coverage of our domain theories to fit all of the mechanics ranking problems in [Hewitt, 2010]. Then, the utility of these capabilities will be explored by incorporating them into the sketch-based educational software systems we are building. Ultimately, we would like to extend coverage to all of conceptual physics, to create a Socratic tutor that can help students understand physics, using a combination of interacting sketching and language.

Acknowledgments

This work was supported by the Spatial Intelligence Learning Center (SILC), an NSF Science of Learning Center.

References

- [Chaudhri *et al.*, 2009] Chaudhri, V., Clark, P., Mishra, S., Pacheco, J., Spaulding, A., & Tien, J. 2009. AURA: Capturing knowledge and answering questions on sci-

- ence textbooks. Technical Report, SRI International, Menlo Park, USA
- [Cohn *et al.*, 1997] Cohn, A. G., Bennet, B., Gooday, J., and Gotts, N. M. 1997. Qualitative Spatial Representation and Reasoning with the Region Connection Calculus. *GeoInformatica* 1: 275-316.
- [Falkenhainer *et al.*, 1989] Brian Falkenhainer, Kenneth D. Forbus, and Dedre Gentner. The structure-mapping engine: Algorithm and examples. *Artificial Intelligence* 41(1): 1-63. 1989.
- [Falkenhainer & Forbus, 1991] Falkenhainer, B. and Forbus, K. (1991). Compositional modeling: Finding the right model for the job. *Artificial Intelligence*, 51, 95-143.
- [Forbus 1984] Forbus, K. (1984). Qualitative process theory. *Artificial Intelligence*, 24, 85-168.
- [Forbus *et al.*, 2011] Forbus, K., Usher, J., Lovett, A., Lockwood, K., & Wetzel, J. (2011). CogSketch: Sketch Understanding for Cognitive Science Research and for Education. *Topics in Cognitive Science*, 3: no. doi: 10.1111/j.1756-8765.2011.01149.x
- [Gentner, 1983] Dedre Gentner. Structure-mapping: A theoretical framework for analogy. *Cognitive Science* 7(2): 155-170. 1983.
- [Graesser *et al.*, 2003] Graesser, A.C., Jackson, G.T., Mathews, E.C., Mitchell, H.H., Olney, A., Ventura, M., Chipman, P., Franceschetti, D., Hu, X., Louwerse, M.M., Person, N.K., & TRG (2003). Why/AutoTutor: A test of learning gains from a physics tutor with natural language dialog. In R. Alterman and D. Hirsh (Eds.), *Proceedings of the 25rd Annual Conference of the Cognitive Science Society* (pp. 1-5). Boston, MA: Cognitive Science Society.
- [Hewitt, 2010] Hewitt, P. 2010. *Conceptual Physics*. 11th edition, Pearson-Addison Wesley, Boston.
- [Jee *et al.*, 2009] Jee, B., Gentner, D., Forbus, K., Sageman, B. and Uttal, D. (2009). Drawing on experience: Use of sketching to evaluate knowledge of spatial scientific concepts. In *Proceedings of the 31st Annual Conference of the Cognitive Science Society*. Amsterdam, The Netherlands.
- [Kim, 1993] Kim, H. (1993). Qualitative reasoning about fluids and mechanics. Ph.D. dissertation and ILS Technical Report, Northwestern University.
- [Kindfield, 1992] Kindfield, A. C. H. (1992). Expert Diagrammatic Reasoning in Biology. In *Proceedings of the AAAI Spring Symposium*. Palo Alto, CA.
- [Klenk *et al.*, 2005] Klenk, M., Forbus, K., Tomai, E., Kim, H., and Kyckelhahn, B. (2005). Solving Everyday Physical Reasoning Problems by Analogy using Sketch-*es. Proceedings of 20th National Conference on Artificial Intelligence*. Pittsburgh, PA.
- [Larkin and Simon 1987] Larkin, J. and Simon, H. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 11:65-99.
- [Leonard *et al.*, 1996] Leonard, W. J., Dufresne, R. J. and Mestre, J. P. (1996). Using qualitative problem-solving strategies to highlight the role of conceptual knowledge in solving problems. *American Journal of Physics*, 64(12):1495-1503.
- [Lockwood & Forbus, 2009] Lockwood, K. & Forbus, K. (2009). Multimodal knowledge capture from text and diagrams. *Proceedings of KCAP-2009*.
- [Lovett & Forbus, 2010] Lovett, A. & Forbus, K. 2010. Shape is like space: Modeling shape representation as a set of qualitative spatial relations. In *Proceedings of the AAAI Spring Symposium*. Palo Alto, CA.
- [Nielsen, 1988] Nielsen, P.E. (1988). A qualitative approach to rigid body mechanics. (Tech. Rep. No. UIUCDCS-R-88-1469; UILU-ENG-88-1775). Urbana, Illinois: University of Illinois at Urbana-Champaign, Department of Computer Science.
- [Novak & Bulko 1990] Novak, G.S. & Bulko, G.S. (1992). Understanding natural language with diagrams. *Proceedings of the eight national conference on Artificial Intelligence*. Boston, MA.
- [Rajagopalan, 1995] Rajagopalan, R. 1995. Picture semantics for integrating text and diagram input. *Artificial Intelligence Review* 10(3-4). Special issue on Integration of natural language and vision processing.
- [Rose *et al.*, 2001] Rose, C., Jordan, P., Ringenberg, M., Siler, S., VanLehn, K., & Weinstein, A. 2001. Interactive conceptual tutoring in Atlas-Andes. In J. D. Moore, C. L. Redfield, & W. L. Johnson (Eds.). *Artificial Intelligence in Education* (pp. 256-266). Amsterdam: IOS Press.
- [Weld, 1988] Weld, D. 1988. Comparative Analysis. *Artificial Intelligence*, 36:333-374, October 1988.
- [Wetzel & Forbus, 2008] Wetzel, J. and Forbus, K. (2008). Integrating Open-Domain Sketch Understanding with Qualitative Two-Dimensional Rigid-Body Mechanics. In the *Proceedings of the 22nd International Workshop on Qualitative Reasoning*. Boulder, CO.
- [Wetzel & Forbus, 2010] Wetzel, J. and Forbus, K. 2010. Design Buddy: Providing Feedback for Sketched Multimodal Causal Explanations. In *Proceedings of the 24th International Workshop on Qualitative Reasoning*. Portland, Oregon.
- [Yin *et al.*, 2010] Yin, P., Forbus, K., Usher, J., Sageman, B. & Jee, B. 2010. Sketch Worksheets: A Sketch-based Educational Software System. *Proceedings of the 22nd Annual Conference on Innovative Applications of Artificial Intelligence*.