The Physics of Futures Past:
A response to Sacks and Doyle

Kenneth D. Forbus

Qualitative Reasoning Group
The Institute for the Learning Sciences
Northwestern University
1890 Maple Avenue
Evanston, Illinois, 60201

Internet: Forbus@ils.nwu.edu
Voice: (708) 491-7699

1 Introduction

In [45], Sacks and Doyle claim to "evaluate the success of the qualitative physics enterprise in automating expert reasoning about physical systems." They present their view of the field, complain that this approach is inadequate, and suggest that an alternate approach would be better. Unfortunately, each part of their paper is either incorrect, misguided, or both. That is, (1) their view of qualitative physics, "simulation of processes by qualitative reasoning" (SPQR), is inaccurate and misleading and (2) their account of expertise is wrong. This essay first argues that the aims and progress of qualitative physics, as described in the literature, differ substantially from Sacks and Doyle's description, and second, that a close examination of how experts work suggests that the ideas and techniques developed in the last 15 years in qualitative physics will play a central role in making computers approach the depth and flexibility of expert reasoning about physical systems. Some of the suggestions phrased by Sacks and Doyle as a radical break from current practice, are actually part of the way work in the field goes – hence the title of this essay.

2 SPQR isn't qualitative physics

All quotes, except where otherwise noted, are from [45].

> "Most of this research seeks to capture the ability of experts to predict the behavior of dynamical systems, such as circuits, fluid flows, and mechanisms."
>
> "Although each approach has its own special characteristics, they share a modeling language for physical systems, a representation for behavior, and an analysis method."
>
> "After fifteen years of research, SPQR can analyze successfully only a handful of simple systems, ..."
>
> "... the analyses that appear in current research papers seem, from a mathematical point of view, little better than those of five years ago."
>
> "Concerns with producing causal explanations and with explaining commonsense reasoning may underlie the differences between SPQR and expert knowledge and methods, ... But both these positions are controversial. We will not address either of these issues here in order to focus more clearly on the accepted aims and methods of SPQR."

Sacks and Doyle assert that prediction is the sole goal of qualitative physics, and that simulation is the sole method of analysis. They complain that qualitative physics isn't making progress. Even a cursory glimpse at the literature shows that each of these statements is false. This section reiterates what the goals of qualitative physics actually are, and shows that concerns which Sacks and Doyle attempt to dismiss – such as integration with other forms of knowledge, modeling, causality, and explanation – have been, and continue to be, central to the enterprise[1].

---

[1] In this essay I will mainly draw on my own work and that of my students, since other respondents will, given the size restrictions, no doubt do the same. For most of my points there are many other examples which would do just as well.

2.1 The real goal of qualitative physics

> "More to the point, it [the qualitative physics literature] lacks any examples where experts draw incorrect or incomplete conclusions while SPQR does better on a generalized version."

The goal of qualitative physics is to characterize the kinds of representations and reasoning used in commonsense thinking about the physical world. There are two interrelated motivations for this goal. One motivation is understanding human reasoning and learning more generally (c.f. [15,29,47]). Since the person on the street rarely knows differential equations (and, as we show below, neither do many experts!), clearly something else must be going on, and finding out how such reasoning works is a legitimate and important intellectual enterprise. (Sacks and Doyle presumably have no argument with this aspect of the field, since their arguments concern expert reasoning.)

The other motivation for qualitative physics is to capture the *tacit knowledge* used by scientists and engineers. Scientist and engineers know much about the physical world before they start their professional training. The claim is that if one closely examines how they reason, one finds that their commmonsense theories of the physical world (often refined as a consequence of their schooling) provides a substrate for their professional skills. That is, a qualitative physics can be used to organize other kinds of knowledge, including quantitative knowledge[2]. This argument has been made repeatedly in the literature, and many of Sacks and Doyle's complaints are based on denying this premise. We return to the issue of how experts reason in Section 3.

Qualitative physics is intended to play a supporting role in some engineering tasks and a more central role in others. Tasks where qualitative knowledge appears central include conceptual design, recognizing function from structure, troubleshooting, monitoring, and instruction. These tasks have inspired work on several styles of reasoning, including measurement interpretation [24,11,33,14], comparative analysis [50], diagnosis [7], learning [16,42,47], and explanation [28,31]. These broad goals have been articulated from the very beginning[3]. Thus even a cursory reading of the literature does

---

[2] To quote [22], p 91, "A purely qualitative theory cannot hope to capture the full scope of human reasoning about physical domains. However, by defining a basic theory using qualitative representations, we can later add theories involving more precise information ... to allow more precise conclusions. ... In this way we can add theories onto a common base that capture more sophisticated reasoning, such as an engineer uses when estimating circuit parameters or stresses on a bridge."

[3] For example, the catalog of styles of reasoning from [22] included determining activity, prediction, postdiction, skeptical analysis, measurement interpretation, experiment planning, and causal reasoning.

not support the assertion of Sacks and Doyle that prediction is or was ever the only important goal.

Simulation, ranging from purely qualitative to various mixtures of qualitative and numerical information, has received much attention. Partly this is because, as Sacks and Doyle rightly point out, simulation can be a useful component in some engineering tasks. It is also because evaluating simulations is relatively easy — constructing an entire design or tutoring system, for instance, requires building many additional parts. And some of the attention is probably due to a propensity to build on last year's conference paper, or a conserative bias. However, Sacks and Doyle never cite any specific paper as claiming that its goal was to out-perform the predictive accuracy of numerical simulation with purely qualitative techniques. This is probably because there isn't one. Qualitative simulation has different goals than predicting detailed behavior, as Section 2.5 describes.

2.2 Myth: Only purely qualitative need apply

> "Qualitative physics eschews numerical analysis, arguing that it provides only reams of numbers, not qualitative information, and that it is prohibitively expensive and unreliable for analyzing realistic systems"
>
> "But SPQR is a theory of expert reasoning that makes use of no knowledge or methods that are not perfectly intelligible to the educated layman, and its extensions draw on only a minor portion of the knowledge visibly used by experts in their reasoning."
>
> "The current practice, with its implication that mathematics has never addressed the problems of reasoning in any significant way, and with the concomitant "not invented here" requirement that one must abandon established concepts and methods to automate expert reasoning ... should be abandoned in favor of a more productive spirit of cooperation and building on past discoveries."

One wonders whether Sacks and Doyle have actually read any of the literature they cite. Qualitative physics has from the beginning been concerned with exploiting other kinds of knowledge. There are many examples in the literature, ranging from algebraic manipulation [52], teleology [10], and configuration space [20]. Even if one remains focused on numerical simulation, as do Sacks and Doyle, there are many examples, such as the use of diagrams to provide geometric boundary conditions for constrained-based numerical simulation [21] and using analytic bounds to prune possible behaviors [36].

It is true that as we discovered that a surprising amount could be done with purely qualitative knowledge, a strong effort was made to find out where qualitative knowledge alone would succeed and where it would fail. This is as it should be — studying factors in isolation is a classic method in science.

Some of the most interesting results in qualitative physics concern the interaction between qualitative and other kinds of knowledge. For instance, in spatial reasoning, the *poverty conjecture* [27,32] argues that there is no purely qualitative, task-independent representation for shape and space. Instead, metric diagrams must be used to compute appropriate qualitative representations for specific purposes. Another example is solving textbook problems. The first work in qualitative physics [8] was an exploration of how qualitative reasoning could be used to guide the construction and solution of equations in simple motion problems as well answering simple questions directly. Later, Skorstad and I showed that qualitative analysis plays a similar, central role in solving engineering thermodynamics problems, with roughly one-third of the equations needed to solve sample problems coming from reasoning using the qualitative model [48]. A third example is the notion of self-explanatory simulation introduced by Falkenhainer and myself [31], which combines the precision of numerical simulation with the explanatory power of qualitative representations. The causal structure derived by the qualitative analysis provides the organization needed to weave equation fragments into simulation code, as well as providing explanation facilities. A fourth example is the whole genre of qualitative analysis of the results of numerical simulations, including the work of Sacks [44], Yip [54], Nishida [40], and Zhao [55].

These cases by no means exhaust the examples from the literature which explore the relationship between qualitative knowledge and other forms of knowledge. However, they should suffice to make the point that qualitative physics does not "eschew" numerical analysis, and that there is no "not invented here requirement" in such research. The only substantiation Sacks and Doyle provide for their complaints is a quotation from one of my papers (Section 3.3 of [45]). An examination of the source indicates that this quotation is taken out of context. The example being discussed is what a robot needs to know in order to make coffee. Do Sacks and Doyle seriously propose that knowing specific differential equations and analyzing them with advanced mathematical techniques is necessary to perform this task? The rest of that section in the original paper argues that similar problems do show up in other kinds of reasoning about the physical world, so Sacks and Doyle may have thought the quotation appropriate. But the rest of that section also highlights both the importance of modeling and how qualitative knowledge interacts with other kinds of knowledge! This flatly contradicts their claims. Let us examine the modeling issue in more detail next.

2.3 Modeling has always been central

> " . . . a handful of simple systems, such as a falling point mass or a U-shaped tube containing liquid."
>
> "qualitative physics research should focus on modeling and automating existing mathematics rather than on inventing analysis tools."

Modeling has been a central theme of qualitative physics since its inception. Many of the theories explore different ontological commitments (e.g., component-centered [9] versus process-centered [22], viewing fluids as contained stuffs versus pieces of stuff [4]). With the exception of `QSIM` [35], no qualitative reasoning system starts with a set of equations as its input. Instead, sets of equations (sometimes qualitative, sometimes quantitative, sometimes both) are derived from some different description. It is hard to see how that is not modeling!

Sacks and Doyle's summary of examples misses the point. For instance, they describe the two containers and three containers examples from my own work as "U-tube" and "W-tube" examples. This description already presumes the results of most of the reasoning that a QP interpreter is designed to capture. Figure 1 shows how the two-container example is typically described to a QP interpreter. Notice that the description concerns the physical structure of the scenario; the scenario does not contain equations, either qualitative or quantitative. Instead, it is described in terms of objects and some relationships that hold between them. The work of modeling, that is, of deriving a description of this system which is useful for some task, is a major portion of the work a QP interpreter performs [25].

A variety of models can be generated for even the simple scenario of Figure 1. Geometric aspects, thermal properties, limitations on phases considered (i.e., liquid, gas, or both), and fault models (e.g., whether or not to worry about paths being blocked) all may or may not be included, depending both on the domain theory used and the task at hand. The set of modeling choices becomes even larger when quantitative concerns are included — for instance, should flows be modeled as laminar or turbulent?

Calling this scenario a "U-tube" is misleading because presumes that many of these choices are made outside the scope of the program's reasoning. For this scenario to be accurately modeled as a U-tube, the fluid path must not be not blocked and there must always be fluid in both containers and

Figure 1: Description of two containers problem

The assertions below show how the two containers problem is described to programs like QPE. Notice that it does not include equations, either qualitative or quantitative. Constructing a model of the scenario, based on knowledge of a domain, is part of the QP interpreter's task, not something that is given as the input. Thus Sacks and Doyle's description of this as a "U-tube" begs many interesting issues. Even reasoning about this simple example requires choosing whether to consider geometry, whether paths can be blocked, whether to include thermal effects, and whether to included gases as well. Figuring out how to apply knowledge of the physical world appropriately has always been a central concern of qualitative physics.

```
state(liquid)          Fluid-Connection(P1,F,G)
substance(water)       Fluid-Connection(P1,G,F)

Container(F)           A[bottom-height(F)] = A[max-height(P1)]
Container(G)           A[bottom-height(G)] = A[max-height(P1)]
Fluid-Path(P1)
```

the path. But arriving at this simplification is part of the modeling process being formalized. For instance, the fact that if there is fluid in both containers then there will always be fluid in both containers is inferred qualitatively, not stipulated. Changing the geometric assumptions, for instance by placing the bottom of one container higher than another, leads to a different set of qualitative behaviors for which the U-tube simplification is inappropriate.

Contrary to Sacks and Doyle's suggestions, there is no pre-existing mathematical formalism that captures this modeling process. But Qualitative Process theory, especially as augmented by the compositional modeling methodology [17,18] can. For instance, Falkenhainer and I have demonstrated that choices of physical models for numerical simulation, such as laminar versus turbulent flow, can be reasoned about using compositional modeling [19,18].

QP theory and compositional modeling are just one example among many of modeling work in the literature. Others include the work of Addanki's group on the Graph of Models framework [2], Weld's work on approximations[49], Liu and Farley's work on integrating multiple ontologies [38]. The only evidence Sacks and Doyle provide for their claim that qualitative physics ignores modeling is a quotation from an introduction in [51]. This seems to be another case where flamboyant polemic is rooted in poor scholarship.

2.3.1 Losing the haystack while looking for the needle

There is a certain circularity in Sacks and Doyle's argument concerning the supposed lack of progress in qualitative physics. They begin by stipulating that accurate prediction is the only goal, and that therefore the complexity of equations which can be handled is the appropriate measure of progress. They proceed to survey the literature by classifying the examples in terms of equations (often presumed by them, not from the work in question), and declare that there has been little progress. Sacks and Doyle state they excluded examples when the papers did not list equations. But this is indulging in a process of self-selection designed to buttress their point. Papers, especially conference papers, rarely provide detailed descriptions of large examples.

There is no reason to believe that counting equations in papers is an appropriate measure of progress. A more reasonable measure would be to look at the complexity of systems being modeled and the kinds of analyses performed upon them. And by these measures, there has been substantial progress. For instance, simple versions of refrigerators [4] and steam propulsion plants [17,48] have been successfully reasoned about, as well as analog electronic circuits of the complexity of operational amplifiers [10]. The breadth of tasks is also wider than they suggest, including the solution of textbook problems, question answering for instruction, and the automatic generation of simulation programs from structural descriptions. Both the complexity of systems modeled and the sophistication of the reasoning with those models is higher than the impression given by Sacks and Doyle.

2.4 Causality and explanation are central

One of the key motivations of qualitative physics is to capture commonsense reasoning and causal explanations. Sacks and Doyle dismiss this aim as "controversial", citing an introduction to a section on causal reasoning in [51]. What Weld and de Kleer actually say is that causality is one of "the most fractious topics", which is correct. But the arguments are not about whether causal reasoning is important: they concern how best to formalize it. There are indeed many differences between the various approaches to qualitative physics; the fact that causality stirs such arguments is a good indication that causality is of central importance.

Although many of those motivated by "hard-core" engineering problem solving who do qualitative physics research are loath to admit it, cognitive fidelity is a crucial criterion. Explanation and causal reasoning, for instance, are judged in part by whether or not they make sense to human listeners. Qualitative physics provides a new opportunity to build systems whose knowledge has a good conceptual impedance match to that of the humans it interacts with. This is important for making programs that are good team players, rather than oracles [39]. It is also essential in tutoring — qualitative physics provides a formalism for encoding human mental models [15], so that a system can represent the kinds of knowledge students are likely to have. And, in Section 2.6, we see that causal reasoning is useful for control and diagnosis as well.

2.5 What is an analysis, anyway?

> "whatever one's position on causality, the limitations of SPQR reduce its utility for producing causal explanations, since it cannot explain systems it cannot analyze."

Sacks and Doyle stipulate that an analysis has failed if it contains ambiguity. If the goal of qualitative physics was, as they assert, to perform exactly the same tasks as traditional numerical simulation this would not be unreasonable. But, as we have seen above, this has never been the goal. The existence of ambiguity in qualitative reasoning was recognized from the beginning. What Sacks and Doyle fail to understand is that, given the complementary role qualitative physics plays with traditional techniques, the right kind of ambiguity is actually desirable.

Think of an envisionment as a grammar for behaviors. In an accurate envisionment every physically realizable behavior of the system corresponds to a path through it, just as every legal sentence corresponds to some traversal of the nodes and arcs in a grammar. The fact that a particular grammar for English may not, given a few words of a sentence, always predict the next word uniquely does not mean that the grammar has failed to analyze a sentence. There may be several legal possible ways to complete the sentence. And so it is with envisionments. Even very little information about a system can suffice to provide a characterization of the kinds of behaviors (i.e., qualitative states) that may ensue.

The ability of envisionments to capture possible classes of behaviors can be exploited in several ways. In measurement interpretation, for instance, numerical descriptions of behavior are explained in terms of qualitative states by a computation similar to parsing [24], although extra subtlety is needed to handle duration information, sensor failures, and certainty information [12]. Such information also should be useful in design. Given a abstract structural description of a system, an envisionment can identify the classes of dangerous states which may or may not arise, depending on the details of the design.

Choosing the properties and parameters of the artifact so that these behaviors are not actually possible would become one of the goals of the design. Incorporating the possible actions the operators of the artifact might take can provide a means for verifying the effectiveness and safety of operating procedures [30,13]. The ability to examine a space of possible behaviors before numerical parameters (or even detailed structural descriptions) are chosen can allow potential mishaps to be identified in the early, conceptual stages of design, when errors are less costly. Even if prediction is one's goal, envisionments can still useful. If there is a unique behavior predicted in a correct envisionment, then whatever behavior occurs must satisfy that qualitative behavior. If there are several possible behaviors, the explicit branching in the representation makes it easy to see where additional knowledge is required to resolve the ambiguity.

2.6 When qualitative is better

> "Qualitative physics offers two arguments that specific equations cannot adequately model physical systems whose exact workings are unknown or extremely complex: specific equations must incorporate unwarranted assumptions about the system, hence may yield incorrect inferences; and specific equations may yield no inferences due to limitations on the reasoning or computational abilities of the analyst. Although both problems can occur, the qualitative physics literature contains only anecdotal evidence that they occur in practice."

Few professions are rewarded for documenting failures, and science and engineering are not among them. So it should surprise no one that documentation of limitations in engineering practice is hard to find. But there are plenty of obvious clues that numerical simulation alone does not constitute engineering nirvana. If FORTRAN numerical programs sufficed for diagnosis and monitoring, why does the engineering community show so much interest in expert systems? The problems are there, if you talk to scientists and engineers. And researchers in qualitative physics do talk to engineers, and have documented such cases, contrary to Sacks and Doyle's allegation. This section briefly mentions two examples.

The work of LeClair's group on Qualitative Process Automation (QPA) [37] demonstrates that qualitative techniques can do a better job than statistical process control techniques. In curing composite parts the temperature of the furnace needs to be kept relatively low while the part is outgasing.

But keeping the furnace low means the parts take longer to cook, reducing productivity. So the best strategy is to keep the heat low while outgasing is occurring, and once it is finished increase the heat to finish cooking the part more quickly. Statistical process control methods use a combination of analytic models and empirical tests to figure out an optimal pattern of high/low cooking times. QPA incorporates a qualitative description of behavior into the controller, allowing it to detect the change in qualitative regime and control the furnace accordingly. The use of qualitative distinctions in supervisory control provided both faster cooking times and higher yield rates than traditional techniques.

The work of Abbott's group on operative diagnosis in civilian aviation [1,46] is another example of where qualitative models have proven to be useful. Few accurate numerical fault models of failures in turbofan aircraft engines have been developed. There are three reasons: (1) the phenomena is hard to model, (2) there are many classes of failures and their possible effects, and so no complete catalog of faults currently exists, and (3) the validation of such models would be fantastically expensive. (Even if such numerical fault models existed they would be irrelevant, as Section 3.3 explains.) Instead, Abbott's `FAULTFINDER` system uses a qualitative engine model to represent fault hypotheses and their consequences. While field tests still lay in the future, evaluation of the system against existing incident reports has been extremely encouraging [46]. Importantly, the use of qualitative models allows their system to produce easily interpretable summaries of system status, which should be invaluable under flight conditions.

What these efforts have in common is that (a) they are driven by real-world tasks and (b) they are using qualitative physics to do things for which traditional techniques have been irrelevant. As with other work in qualitative physics, they are not shy about using traditional techniques where appropriate: `QPA` includes numerical parameters within its qualitative decomposition of curing, and `FAULTFINDER`'s diagnosis process is driven by discreprencies detected between engine operation and a simulation of a normally working engine[4]. Again, we have more evidence that Sacks and Doyle's view of qualitative physics bears little relation to the reality of it.

---

[4] Models of normal operating regimes do generally exist, a fact which is also routinely exploited in building training simulators. However, the lack of good fault models often limits the kinds of training they can deliver, because the trainer will halt when it goes beyond the limits of its simulation model.

2.7 Has Dynamics done it all?

> "some basic mathematical knowledge and methods have already been automated with considerable success."

Sacks and Doyle are willing to concede that there has been some good work in qualitative physics (e.g., [44]). Since they proclaim the superiority of numerical simulation for all engineering tasks, it is interesting to ask just how much progress has actually been made in the qualitative interpretation of numerical simulations. The record suggests that work in this part of qualitative physics actually lags behind work in other areas, in terms of the kinds of systems it can handle with respect to its goals.

Until recently, all programs which perform qualitative interpretation of numerical simulations have been limited to one or two independent state parameters. Other restrictions are sometimes required as well (e.g., non-dissipative maps). But other research in qualitative physics has routinely involved examples which contain more than two independent state parameters, such as abstract models of refrigerators and steam propulsion plants[5] To be sure, the goals have been different: In most analyses of thermodynamic cycles, steady-state conditions are all that matter. Under such constraints envisioning is easy and efficient, yielding only a handful of states which can be discriminated between by taking the intended function of the system into account [48]. But in any case, it seems that the current generation of systems advocated by Sacks and Doyle are actually less capable than systems which do not use numerical simulation for their tasks.

Scaling is always hard, of course, but there is no evidence that the simulate-then-interpret approach will scale better than direct qualitative reasoning, and there are reasons to suspect it will do worse for many tasks. The recent work of Zhao [55] suggests that such techniques can escape the plane, so the limitation of dimensionality may not be serious. However, as the number of independent system parameters grows, the cost of each numerical experiment grows and, more seriously, the number of numerical experiments required to sample the phase space grows exponentially. Can Sacks and Doyle show a numerically-simulate-then-interpret program from the literature that can handle examples like refrigerators and steam plants? (Even if such a program existed, there is still the crucial question of where the model came from in the first place — see Section 3.6.)

---

[5] In fact, even the two-containers example has eight state variables if thermal properties and both liquids and gases are included.

2.8 Assorted Errors

This section catalogs several inaccuracies not already addressed. This list is not exhaustive.

> "Although the superficial details of SPQR most closely resemble QSIM, which we find especially clear and precise, our discussion applies equally to other approaches."

In fact the only effort SPQR comes anywhere near describing is the `QSIM` research effort. QSIM's use of landmark introduction is very useful for representing fine distinctions concerning specific behaviors. But landmark introduction introduces a different kind of ambiguity than found in envisionments, and has in many ways been more problematic [26]. This has caused Kuipers' group to spend much of their efforts addressing this problem, and these papers may be the source of many of Sacks and Doyle's complaints.

Looking at the literature on Confluences [9] or Qualitative Process theory yields an entirely different picture concerning the relative importance of simulation. For example, confluences have been used extensively in recognizing function from structure (c.f. [10]) and for diagnosis (c.f. [7]). Qualitative Process theory has been used extensively in machine learning and modeling scientific discovery [47]. Both have been used for monitoring [24,11,33]. All of these tasks use the explanatory properties of qualitative physics, and none of them could be carried out by numerical simulation alone.

> "Crawford et. al. (1990) prove that `QSIM` subsumes the dynamics module of QP theory by implementing a translator from QP theory to `QSIM`."

While `QPC` [6] is an interesting system, Sacks and Doyle are drawing an incorrect conclusion from it. QP theory describes how to perform temporal projection involving objects with finite temporal extent [22,23]. Nowhere in the formalization of `QSIM` is there a description of how to handle continuity when objects have a finite temporal extent. (In fact, `QPC` contains techniques *external to* `QSIM`, drawn from QP theory, to perform such inferences so it can accurately implement QP theory.) Therefore `QSIM`'s account of dynamics cannot subsume that of QP theory. And since QP theory does not include landmark introduction (although such an extension would be straightforward), QP theory does not subsume `QSIM`[6] So in their current form the dynamics

---

[6] These differences are mentioned because they should be obvious to anyone familiar with the literature; there are other less important differences as well.

module of QP theory and `QSIM` theory overlap, but neither subsumes the other.

> "After 15 years of research, SPQR can analyze successfully only a handful of simple systems, such as a falling point mass or a U-shaped tube containing liquid. It struggles with linear oscillatory systems, such as simple springs, and fails completely on many non-linear oscillators."

The omission of examples like steam plants and refrigerators, and the inappropriateness of Sacks and Doyle's sampling methodology, has already been mentioned. But the non-inclusion of mechanical systems, such as clocks [20,41] is also puzzling, since Sacks has recently begun to do research in qualitative kinematics also. The most glaring omission, though, is `QUAL`, which was used to analyze hundreds of distinct circuits [10]. Most standard electronic circuit components have well-known analytic models. Sacks and Doyle have been willing to introduce equations in describing the research of others, so these examples would seem natural for them to include in their evaluation.

> "The SPQR states of a system translate to rectangular regions (including degenerate rectangles, such as points and lines) in phase space."

This presumes limit points must be constant, which is false. See [22] for counterexamples.

3 Sacks and Doyle aren't experts on experts

The picture of expertise Sacks and Doyle paint is crucial to their argument: If all experts are numerical simulation jockeys (or should be), and nothing less than the details of full numerical simulation can ever be useful, then indeed the claims of qualitative physics concerning its relevance to expert reasoning would be false. However, their characterization of experts is at best idiosyncratic. The only class of experts which comes close to fitting the Sacks and Doyle account of expertise are dynamicists, preferably a professor in an engineering department, who is mathematically sophisticated, solves problems which do not require immediate action, and has access to reasonably powerful computing equipment. (And even for academic dynamicists there are problems with Sacks and Doyle's account (see Section 3.6).)

There are many dimensions to expertise, and focusing on the use of advanced mathematics provides an incomplete, and hence inaccurate, account of the phenomena. This section demonstrates that Sacks and Doyle's account of experts is incorrect, and highlights the importance of qualitative representations and reasoning in modeling expertise.

3.1 Many experts don't know differential equations

> "Experts reason about dynamical systems by formulating and analyzing differential equations that capture the properties of interest and abstract away irrelevant details."
>
> "Experts reason about dynamical systems with advanced mathematics, physics and other knowledge."

Some experts indeed do work like this. But many don't. Consider a television repairer, auto mechanic, or power plant operator. All of these jobs involve reasoning about complex dynamical systems. All of them involve substantial expertise. Yet anyone who has observed these experts in action knows that what they do bears little resemblance to the Sacks and Doyle account of expertise. I have never seen an auto mechanic, even very good ones, use a numerical simulation in troubleshooting my car. Nor do power plant operators respond to alarms by pulling out a notebook and attempting to construct an analytic or numerical model of the current problem.

One can be an expert auto mechanic with a high-school education and without knowing what a differential equation is. The same holds true in many monitoring and troubleshooting jobs (c.f. [53]). Students at the U.S. Navy Surface Warfare Officer's School are college graduates without engineering degrees, and typically have never taken a differential equations course. They are taught a substantial amount of the physics and operation of steam plants, including ample numerical information and some basic equations. But they never learn the fundamentals of differential equations nor how to solve them, either analytically or numerically.

3.2 Many experts couldn't use differential equations if they had them

> "experts need to know the asymptotic behavior of a dynamical system: the stable steady states the sets of solutions that converge to each steady state and the sensitivity of these properties to perturbations in the equations. This information provides

> a qualitative understanding of the system and sets the stage for further analysis ..."

I am writing this on board a commercial aircraft, and I sincerely hope that (a) my pilot is an expert and (b) that he or she is *not* thinking right now about the asymptotic behavior of approximations to the flight dynamics of this 747! Even if my pilot knows the differential equations of flight, they would be fairly useless to him or her in this operational context. In monitoring and control tasks, actions are associated with the qualitative regimes of behavior they are appropriate for[7].

Suppose, heaven forbid, something goes wrong with the plane. There are not very many sensors to provide detailed information about the internal state of the aircraft's systems. Even if there were, there is already so much information competing for a pilot's attention that overload is a serious problem. These are not the circumstances in which one wants to be developing new differential equation models of aircraft subsystems! There are simply too many ways things can go wrong, and no good ways to validate the kinds of detailed models that Sacks and Doyle suggest are the basis of expertise.

This is not an isolated situation with aircraft operation and design. For instance, existing simulation systems for chemical process plants focus on steady-state behaviors and contain few models of transient behaviors or fault conditions. Training simulators of all kinds, from flight simulators to propulsion plant trainers, only model a subset of the kinds of behaviors their users may encounter in real life due to the difficulty of coming up with reasonable numerical models of system failures. Numerical simulation, while very useful for some tasks, is not the panacea suggested by Sacks and Doyle.

3.3 Would dynamicists be good pilots?

Suppose there were a complete set of differential equation models for every possible failure and fault in an aircraft. Would they be useful? Let's assume a loss of hydraulic pressure occurs. Somehow the proper fault model would have to be identified. If the only level of description allowed is that of concrete, specific equations this task becomes more difficult. The ability of qualitative representations to capture whole classes of behavior simplifies the system identification problem.

Let's assume that we have found the right differential equation model (or simulation program) for the current fault, somehow. What do we want to

---

[7] Of course, even though judging what action to take may be based on qualitative distinctions, the parameters of the action sometimes are computed by machines using solutions to differential equations.

do with it? What a pilot needs to know is how the current problem will affect the safety and operation of the aircraft. This means the model should be able to predict what other failures may occur as a consequence of the initial failure, for example. It also requires interpreting the results of the simulation in terms of the intended functioning and operational goals of the aircraft. The numerical simulation alone does not provide this. Qualitative descriptions are needed.

Information about the structure, function, and behavior of the physical system can be connected through the medium of the representations developed by qualitative physics research. Certain classes of states or properties can be described as dangerous (or at least inconsistent with normal operations), and many operational goals can be described in qualitative terms. Techniques like action-augmented envisionments [30] provide a model for some aspects of the reasoning pilots do when they figure out "work-arounds" for system failures. Formalizing these notions is an important part of qualitative physics. But such concerns appear nowhere in Sacks and Doyle's account, even though they are clearly important to automating engineering problem solving.

Precise models of fault conditions, if they existed and if they could be applied appropriately, might provide some valuable information in an operational context. For instance, if one could estimate the size of a fuel leak accurately better predictions of expected range could be made. However, it is unrealistic to expect such precise information is always available. Sensors are limited and prone to inaccuracies and failures, precise characteristics change as parts wear, etc. And making detailed but inaccurate predictions available to the pilot is at best distracting, and possibly quite dangerous.

What do pilots actually do? Abbott found that pilots base their strategies on qualitative models of faults [1]. Interestingly, Abbott found that pilots often have two kinds of strategies, depending on whether or not they have isolated what is wrong with a subsystem or just know that a particular subsystem is broken somehow. Notice that modeling a system as just "abnormal", as opposed to specifying how it is broken (what Abbott calls the *status abstraction*) is even less precise than the qualitative models deplored by Sacks and Doyle. Yet pilots find this abstraction useful.

Thus we see that, even if we had the kind of detailed models Sacks and Doyle stipulate as the foundation for expertise, we have captured only part of what expert pilots needs to know. And since `FAULTFINDER` demonstrates that it is possible to derive and track fault hypotheses via qualitative reasoning [1], at best the numerical simulation proposed by Sacks and Doyle as a required step in the interpretation process is unnecessary. Experience suggests that this conclusion is reasonable for troubleshooters and plant operators of all kinds.

Sacks and Doyle may not choose to regard pilots, auto mechanics, or even most scientists and engineers as experts. That is certainly their choice. But then they are drawing the line between expert and non-expert much more narrowly than most everyone else, including the qualitative physics community. (Drawn narrowly enough, of course, it becomes tautology: unless someone uses differential equations, they cannot be a "real" expert. This is a fallacy, since expertise should be judged by performance not by method.) Most experts, under the usual definition of the term, do not "have access to fast computers to do numerical simulation" in the course of their daily lives. Yet somehow they manage to get their work done and keep things running safely and smoothly without such advantages.

Perhaps it would indeed be useful to arm auto mechanics with finite-element analysis programs. But that is a different argument. The fact that many human experts, as the term is used normally and in the qualitative physics literature, perform without such knowledge is a strong existence proof that heavy-duty mathematical analysis is not a necessary component of expertise. This fact indeed was one of the major motivations for undertaking the qualitative physics enterprise in the first place, as the literature attests.

3.4 Engineers need more than differential equations

Suppose we limit our focus to engineers, and specifically, those engineers working off-line, in contexts where they do not need to respond immediately. Do they, as Sacks and Doyle suggest, spend their time "revising specific equations until predications derived from these equations become accurate enough for their needs."? Suppose they do. How does the engineer know an equation is accurate enough? The idea of approximation implies some standard for comparison. If the engineer's expectations were sufficiently accurate already, there would be no need for a simulation. So there must be either (a) some external standard of comparison, such as empirical data, or (b) some expectations that are expressed with less precision. Perhaps even with so little precision that there is, well, perhaps some *ambiguity* in the engineer's expectations?

A claim of qualitative physics is that one reason an engineer may choose to do a particular simulation is that she did a qualitative analysis of the system, albeit informally, first. This qualitative analysis framed the questions, if any, which needed to be resolved by a more detailed analysis. The qualitative analysis proposes classes of alternatives, other knowledge (including numerical) disposes of them as necessary. The qualitative summary of the detailed analysis is required to answer the initially posed questions. Notice that the Sacks and Doyle account provides no reason for the engineer to summarize the results of an analysis qualitatively, whereas the standard account of qualitative physics does!

It should be clear that there is more to engineering knowledge than differential equations. This shows up even in Sacks and Doyle's account, hidden in words like "domain knowledge" and "experience." Qualitative physics has taken on the responsibility of figuring out what this knowledge looks like. The version advocated by Sacks and Doyle ignores this issue.

> "And whether or not the man on the street infers that "what goes up must come down" by reasoning with general model, the aeronautical engineer analyzes exact aircraft models by advanced mathematics and by extensive numerics."

It is certainly true that mathematical modeling and numerical simulation is part of what aeronautical engineers do. But it is wishful thinking to believe that numerical simulation is all that engineers do. Numerical simulation occurs relatively late in the design process. An important early stage is called *conceptual design*. Conceptual design is what many think of as "back of the envelope" design: the basic idea(s) of a design are sketched out and looked over to see if they are plausible. Conceptual design occurs long before there are detailed structural descriptions and before numerical parameters are chosen, so numerical simulation is not applicable to it. But many serious and expensive errors are made during conceptual design, and by the time numerical simulation can be used to detect the problems, much effort has been wasted[3]. Many engineers are interested in qualitative physics precisely because it could give them leverage for improving conceptual design.

3.5 Mathematical Macho and the Sociology of Engineering

Sacks and Doyle's misconceptions may stem from an interesting property of the sociology of engineering. It is quite common in the culture of engineering to ignore or devalue anything that isn't an equation. This prejudice is taught early: Students are taught to prize analytic models, because they are easy to manipulate. The formulation of models is viewed as something of a black art, part of the "intuition" of an engineer, only gained through bitter experience.

---

[3] Personal communication, Prof. Steven Lu, Department of Mechanical Engineering, University of Illinois at Urbana-Champaign.

This prejudice is why the term *tacit knowledge* arose in qualitative physics. It is much easier to teach equations, since mathematics provides a crisp representation language. If we can provide a similarly crisp language for the rest of what engineers know, then these other kinds of knowledge can be communicated and reasoned about more easily both by people and machines.

How much of engineering knowledge is something other than equations? Sacks and Doyle, by their call to focus exclusively on analytic and numerical models, presume that non-equational knowledge is virtually non-existent, or trivial. But if one examines engineering texts, one sees a different picture.

If the Sacks/Doyle account of engineers were correct the descriptions of important physical phenomena in handbooks would be dominated by the appropriate equations. Yet this is often not the case. Consider boiling, clearly an important phenomena in engineering thermodynamics. The article on boiling in the Van Nostrand Encyclopedia for Science and Technology [5] does not include a single analytic approximation for boiling. Why? Boiling is quite complex, and to model it accurately requires identifying which physical conditions occur in the particular scenario being modeled [3]. There are two important points here: (1) To even select the appropriate approximation, the situation must be analyzed in other, simpler terms first. Nothing in the Sacks and Doyle programme for qualitative physics would explain how this analysis works, but the standard programme of qualitative physics does. (2) A significant portion of engineering knowledge is currently stated informally, in natural language rather than equations. To automate engineering problem solving requires formalizing this knowledge, too.

Modeling in engineering is currently considered an art. Qualitative physics seeks to make it into more of a science, by building formal languages that capture the (often implicit) conditions for using quantitative knowledge. Such conditions are often described in a mixture of qualitative and quantitative terms. For instance, a laminar model for fluid flow is appropriate when the Reynolds number is less than $2,000$ [3]. The wide use of such inequality conditions to determine applicability of models was one of the motivations for basing a qualitative description of number on ordinal relationships in QP theory.

But the role of qualitative representations in modeling goes deeper than this. The very idea that a fluid flow is occurring is a qualitative fact. In analyzing an existing system it might be directly observed, in designing a system it might be something the engineer intends to occur. But in any case, the concept of a fluid flow is a prerequisite to using specific models of fluid flow. The very idea of an approximation suggests that there is indeed some thing being approximated. By representing these abstract concepts explicitly, qualitative physics (as practiced) provides an organizational scheme for other kinds of knowledge.

3.6 Where do models come from? And what are they for?

The issue of modeling reveals a serious flaw in Sacks and Doyle's proposed programme for qualitative physics. Sacks and Doyle presume that efficient, accurate numerical simulation models always exist for phenomena of interest. But where do these models come from?

There is almost never one single model for a complex physical system. A system can be modeled at varying levels of structural detail and at varying levels of accuracy in the phenomena involved. This is also true for numerical simulation models. Each simulation is fragmentary and partial, tuned towards a specific purpose. For instance, a common practice in aircraft training simulators is to model an engine via interpolation from manufacturer's tables of steady-state operating conditions [43]. Such low-resolution simulators are fine for updating instrument displays realistically. But for analysis of a turbofan engine during design, more detailed thermodynamic modeling of individual components is required. Validating such models is an expensive task, often requiring substantial time using test aircraft. This is typical in engineering modeling. In engineering thermodynamics, for instance, good analytic models of advanced power and refrigeration cycles are still published as research journal articles. And, of course, there are often problems in modeling anything but a few behavioral regimes (Section 3.2).

Since no single model is adequate for all purposes, part of formalizing engineering problem solving must be the development of explicit representations which allow the integration of results from different models. Qualitative physics does this by providing an abstract description of a system and its behavior which can be used to assimilate the results of more detailed analyses. This was demonstrated long ago for the simple domain of $2D$ motion [21], and the ideas needed to do this for more complex domains (e.g., engineering thermodynamics, electronics, and mechanics) are coming together today (c.f. [2,18,31,49]). The roles simulation (qualitative and quantitative) plays in engineering tasks is one of the most important problems being addressed by qualitative physics. Yet Sacks and Doyle do not provide any motivation for their qualitative interpretations of simulations, save their assertion that that is what experts do. It would be interesting to see if they can provide such motivation without using the concepts already formalized in qualitative physics.

3.7 Scientists are more than just simulation jockeys

So far our examples are from engineering. What about scientists? Differential equations indeed play a major role in many scientific disciplines. But there are also many fields of science where differential equations have not been all that useful in modeling – evolutionary biology comes to mind, or psychology. For instance, most cellular biologists do not hypothesize and revise specific equations. They engage in very sophisticated reasoning about for instance biochemical pathways, during the course of which they make hypothesis about the structure and function of molecular and cellular systems. Often equations can be found as part of their theorizing, but in service of other representations. And even in those fields where differential equations are most widely used, one still finds that scientific activity encompasses more than simply writing and using equations. Anyone who reads history of science in more than a cursory manner will discover this, and it is still true today. For example, the division of labor in physics between theorists and experimentalists is well-known. In vision research, some scientists do mathematical modeling, others do neurophysiological experiments, others do psychophysical experiments, while others build computational models. Each specialist is essential to uncovering the truth. All of them regard themselves as scientists, and are so regarded by their communities. Yet at most two of them would be called scientists, by the Sacks and Doyle caricature of science.

3.8 Real Experts do eat quiche

To sum up, Sacks and Doyle's vision of expertise has little to do with reality. We daily put our lives and property in the hands of experts who do not know — and do not need to know — differential equations, numerical simulation, or any of the other advanced mathematical techniques that Sacks and Doyle claim are "essential" to expert reasoning. Their account seems in part to be the product of a sampling bias. It is very easy to overgeneralize, especially if one forms opinions by talking to just a few informants.

4 Six months in the laboratory can save an hour in the library. But which library?

Where does this leave us? Every one of Sacks and Doyle's observations, as summarized in Section 5 of their paper, is wrong:

> "The accepted approach, which we have called SPQR, has successfully analyzed only the simplest systems, while routine expert methods succeed on far more complex examples"

In actuality, Sacks and Doyle's SPQR is a caricature, drawn only by ignoring much of the literature (Section 2). Their notion of what an analysis is is misguided (Section 2.5), and their accounting of what qualitative physics has done is flatly wrong (Section 2.3).

> "Virtually all of the systems analyzed in the literature reduce to just three equations."

Section 2.3 demonstrates that this is not the case.

> "SPQR equations are far too general for practical use. Experts instead hypothesize and revise specific equations until they obtain equations of adequate accuracy"

Section 2.6 shows that qualitative representations can be better, in practice, than traditional techniques. Section 3 shows that experts work quite differently than Sacks and Doyle's account of them suggests.

> "Experts focus on asymptotic behavior, while SPQR focuses on transient behavior"

Neither half of this statement is correct; see Sections 3.2 and 2.1.

> "Experts derive the behavior of dynamical systems with deep mathematics and extensive numerical analysis, whereas SPQR uses little of either."

Again untrue; see Sections 3.1 and 2.2.

Sacks and Doyle appear to believe that qualitative physics exists in a vacuum, ignorant of modern mathematics and engineering. Nothing could be further from the truth. Many of those doing qualitative physics research are in fact engineers. Some are practicing engineers whose career path took them back into research (e.g., Roy Leitch). Others are faculty in electrical, mechanical, chemical, or civil engineering departments (e.g., Lyle Ungar). They have been there from almost the beginning, and their numbers are growing. They participate because they see themselves as helping to develop a body of ideas which complements their traditional techniques, and thus can provide new capabilities.

It is Sacks and Doyle who have chosen to live in a vacuum. Consider the following statement:

> "... mathematics provides a rich, well developed store of qualitative concepts and results of proven utility for reasoning about physical systems in all their aspects not just their dynamics"

Has mathematics provided a theory of monitoring and control? In part, yes. There are theories of observability and controllability, and a rich body of work on control theory. But there are still human operators inside power plants and process plants. Many open questions persist, and techniques from AI are being brought to bear on these problems as often as traditional mathematics these days. And of course there are the operational issues which lie outside the traditional boundaries of control theory, but are becoming acute as physical systems become more complex, such as designing controllers that are understandable to humans interacting with them. So, contrary to Sacks and Doyle's assertion, we can't just look up the answer in a math textbook. Is there a mathematical theory of diagnosis? No. Is there a mathematical theory of design? No. Has the modeling process, of moving from a schematic or blueprint or device to a set of equations (either qualitative or quantitative) been reduced to a mathematical theory? No. Mathematics is of course extremely useful in science and engineering, but Sacks and Doyle's mathematics machismo is naive.

Sacks and Doyle might argue that while one cannot look up everything in a textbook today, we should wait for mathematicians to deliver the answers. What qualitative physics already does, which is to embrace mathematics — *when relevant and appropriate* — seems to be a better strategy. When engineering problem solving research has stabilized to the point where such textbooks can be written, those textbooks will include a considerable contribution from qualitative physics. Consider model-building. Qualitative physics brings the tools of logic to bear on the problem. Without modeling languages that include quantification, for example, it is impossible to build explicit domain theories which can be instantiated as needed to analyze a broad class of situations.

The motivations for their essay are somewhat unclear. Sacks, at least, has done excellent research in the area. While the approach he favors has not been the dominant one in the qualitative physics community, it has always been an important part of it. Papers on computing qualitative summaries from numerical simulations routinely appear in both the qualitative physics workshops and in qualitative physics sessions at AAAI and IJCAI. The Sacks and Doyle essay reflects a lack of knowledge (or understanding) of the qualitative physics literature. This would be bad enough if their paper were an attempted technical contribution. But claiming to evaluate an area implies a responsibility to understand it first. Sacks and Doyle simply failed to do their homework.

Qualitative physics as an enterprise is only beginning. We have yet to uncover all of the organizational principles that allow human beings to reason with huge domain theories that incorporate multiple perspectives, ontologies, and granularities. We have explored only a few points in the space-time tradeoffs for qualitative simulation, let alone the issues involved in other styles of reasoning. Some of the most important work will come when we refocus our attention on psychological modeling, using the information processing constraints uncovered in the development of performance systems to provide new light on human cognition. To make progress, qualitative physics must maintain its broad vision and multi-faceted approach. Qualitative physics does not need to "cast off the fetters of currently accepted methods", as Sacks and Doyle suggest. Rather, it is Sacks and Doyle who should take off their blinders and discover that the qualitative physics they seek is already here, around them.

5 Acknowledgements

This essay benefited from extensive discussions with Dedre Gentner and Dan Weld, as well as discussions with Brian Falkenhainer, Ben Kuipers, Jimi Crawford, Johan de Kleer, Pat Hayes, Jane Malin, John Collins, Ted Kitzmiller, Kathy Abbott, and Brian Williams.

References

[1] Abbott, K. "Robust operative diagnosis as problem-solving in a hypothesis space", Proceedings of AAAI-88, August, 1988.

[2] Addanki, S, Cremonini, R., and Penberthy, J.S. "Reasoning about assumptions in graphs of models", Proceedings of IJCAI-89, August, 1989.

[3] Avallone, E. and Baumeister, T. *Marks' Standard Handbook for Mechanical Engineers*, Ninth Edition, McGraw-Hill Book Company, New York,

[4] Collins, J. and Forbus, K. "Reasoning about fluids via molecular collections", Proceedings of AAAI-87, Seattle, Washington, July, 1987.

[5] Considine, D. and Considine, G. (Eds.) —em Van Nostrand's Scientific Encyclopedia, 6th edition, Van Nostrand Reihold Co., New York, 1983.

[6] Crawford, J. and Farquhar, A. and Kuipers, B. "QPC: A Compiler from Physical Models into Qualitative Differential Equations" Proceedings of AAAI-90, August, 1990.

[7] Dauge, P., Raiman, O., and Deves, P. "Troubleshooting: When modeling is the trouble", Proceedings of AAAI-87, July, 1987.

[8] de Kleer, J. "Multiple Representations of Knowledge in a Mechanics Problem Solver" Proceedings of IJCAI-77.

[9] de Kleer, J. and Brown, J. "A qualitative physics based on confluences", *Artificial Intelligence*, **24**, 1984.

[10] de Kleer, J. "How Circuits Work", *Artificial Intelligence*, **24**, 1984.

[11] Decoste, D. "Dynamic Across-Time Measurement Interpretation", Proceedings of AAAI-90, Boston, August, 1990.

[12] Decoste, D. "Dynamic Across-Time Measurement Interpretation: Maintaining qualitative understandings of physical system behavior", to appear in *Artificial Intelligence*, Fall, 1991.

[13] DeCoste, D. "Toward a qualitative theory of safety control – When and how to panic intelligently", Technical Report No. 14, The Institute for the Learning Sciences, Northwestern University, May, 1991.

[14] Dvorak, D. and Kuipers, B. "Model-based monitoring of dynamic systems." Proceedings of IJCAI-89, August, 1989.

[15] Gentner, D., Stevens, A. (Eds.), *Mental Models*, LEA Associates, Inc., 1983.

[16] Falkenhainer, B. "Learning from physical analogies: A study in analogy and the explanation process", Ph.D. thesis, University of Illinois, Urbana. 1988.

[17] Falkenhainer, B., and Forbus, K. "Setting up large-scale qualitative models", Proceedings of AAAI-88, August, 1988.

[18] Falkenhainer, B. and Forbus, K. "Compositional Modeling: Finding the Right Model for the Job", to appear in *Artificial Intelligence*, Fall, 1991.

[19] Falkenhainer, B. and Forbus, K. "Compositional modeling of physical systems", Papers from the 4th International Workshop on Qualitative Physics, Lugano, Switzerland, 1990. Also to appear in *Recent Progress in Qualitative Physics*, Faltings, B. and Struss, P. (Eds.), MIT Press, 1991.

[20] Faltings, B. "A theory of qualitative kinematics in mechanisms", *Artificial Intelligence*, **41**(1), June, 1990.

[21] Forbus, K. "Spatial and qualitative aspects of reasoning about motion", AAAI-80, Palo Alto, California, August, 1980.

[22] Forbus, K. "Qualitative process theory" *Artificial Intelligence*, **24**, 1984

[23] Forbus, K. "The problem of existence", Proceedings of the Cognitive Science Society, 1985.

[24] Forbus, K. "Interpreting observations of physical systems", *IEEE transactions on systems, man, and cybernetics*, Vol. SMC-17, No. 3, May/June 1987.

[25] Forbus, K. "The Qualitative Process Engine" in Weld, D. and de Kleer, J. (Eds.) *Readings in Qualitative Reasoning about Physical Systems*, Morgan-Kaufmann, 1990.

[26] Forbus, K. "Qualitative physics: Past, present, and future" in *Exploring Artificial Intelligence*, Morgan-Kaufman, 1988.

[27] Forbus, K. "Introducing actions into qualitative simulation", Proceedings of IJCAI-89, Detroit, 1989.

[28] Forbus, K. and Stevens, A. "Using qualitative simulation to generate explanations", Proceedings of the third annual meeting of the Cognitive Science Society, August, 1981.

[29] Forbus, K. and Gentner, D. "Learning physical domains: Towards a theoretical framework" in Michalski, R., Carbonell, J. and Mitchell, T. *Machine Learning: An Artificial Intelligence Approach, Volume 2*, Tioga press, 1986.

[30] Forbus, K. "Introducing actions into qualitative simulation", Proceedings of IJCAI-89, August, 1989.

[31] Forbus, K. and Falkenhainer, B. "Self-explanatory simulations: An integration of qualitative and quantitative knowledge", AAAI-90, August, 1990.

[32] Forbus, K., Nielsen, P. and Faltings, B. "Qualitative Spatial Reasoning: The `CLOCK` Project", to appear in *Artificial Intelligence*, Fall, 1991.

[33] Hamilton, T. "HELIX: An application of Qualitative Physics to diagnostics in advanced helicopters", *International Journal of AI in Engineering*, Volume 3, No. 3, 1988.

[34] Kuipers, B. "Commonsense Reasoning about Causality: Deriving Behavior from Structure" *Artificial Intelligence*, **24**, 1984

[35] Kuipers, B. "Qualitative Simulation" *Artificial Intelligence*, **29**, 1986

[36] Kuipers, B. and Berleant, D., "Using Incomplete Quantitative Knowledge in Qualitative Reasoning", Proceedings of AAAI-88

[37] LeClair, S., Abrams, F., and Matejka, R. "Qualitative Process Automation: Self-directed manufacture of composite materials", *AI EDAM*, 3(2), pp 125-136, 1989.

[38] Liu, Z. and Farley, A. "Shifting ontological perspectives in reasoning about physical systems.", Proceedings of AAAI-90, August, 1990.

[39] Malin, J. et. al., "Making intelligent systems team players: Case studies and Design issues. Volume 1. Human-Computer Interaction Design", Technical report in preparation, 1991.

[40] NISHIDA PAPER FROM AAAI-91

[41] Nielsen, P. "A qualitative approach to mechanical constraint", Proceedings of AAAI-88.

[42] Rajamoney, S. and DeJong, G. "The classification, detection, and handling of imperfect theory problems", Proceedings of IJCAI-87, August, 1987.

[43] Rolfe, J.M. and Staples, K.J. (Eds.) *Flight Simulation*, Cambridge University Press, 1986.

[44] Sacks, E. "A dynamic systems perspective on qualitative simulation", *Artificial Intelligence*, 42(2-3):349–362.

[45] Sacks, E. and Doyle, J. "Prolegomena to Any Future Qualitative Physics", *Computational Intelligence, this issue.*

[46] Schutte, P., Abbott, K., Palmer, M., and Ricks, W. "An evaluation of a real-time fault diagnosis expert system for aircraft applications", Proceedings of the 26th IEEE Conference on Decision and Control, 1987.

[47] Shrager, J. and Langley, P. (Eds.) *Computational models of scientific discovery and theory formation*, Morgan-Kaufmann, Los Altos, 1990.

[48] Skorstad, G. and Forbus, K. "Qualitative and quantitative reasoning about thermodynamics", Proceedings of the Cognitive Science Society, August, 1989.

[49] Weld, D. "Approximation reformulations", Proceedings of AAAI-90, August, 1990.

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

[51] Weld, D. and de Kleer, J. (Eds.) *Readings in Qualitative Reasoning about Physical Systems*, Morgan Kaufman, 1990.

[52] Williams, B. "MINIMA: a Symbolic Approach to Qualitative Algebraic Reasoning", Proceedings of AAAI-88, August, 1988.

[53] Woolf, B., Blegen, D., Jansen, J., and Verloop, A. "Teaching a complex industrial process", Proceedings of AAAI-86, August, 1986

[54] Yip, K. "Extracting qualitative dynamics from numerical experiments" Proceedings of AAAI-87, July, 1987.

[55] Zhao, Feng "Extracting and Representing Qualitative Behaviors of Complex Systems in Phase Spaces", Readings of the Fifth International Workshop on Qualitative Reasoning, May, 1991. (Also to appear in IJCAI-91)