Ideal physical systems

Brian Falkenhainer
Xerox Corporate Research & Technology
Modeling & Simulation Environment Technologies
801-27C, 1350 Jefferson Road, Henrietta, NY 14623

Abstract

Accuracy plays a central role in developing models of continuous physical systems, both in the context of developing a new model to fit observation or approximating an existing model to make analysis faster. The need for simple, yet sufficiently accurate, models pervades engineering analysis, design, and diagnosis tasks. This paper focuses on two issues related to this topic. First, it examines the process by which idealized models are derived. Second, it examines the problem of determining when an idealized model will be sufficiently accurate for a given task in a way that is simple and doesn't overwhelm the benefits of having a simple model. It describes IDEAL, a system which generates idealized versions of a given model and specifies each idealized model's *credibility domain*. This allows valid future use of the model without resorting to more expensive measures such as search or empirical confirmation. The technique is illustrated on an implemented example.

Introduction

Idealizations enable construction of comprehensible and tractable models of physical phenomena by ignoring insignificant influences on behavior. Idealized models pervade engineering textbooks. Examples include frictionless motion, rigid bodies, as well as entire disciplines like the mechanics of materials. Because idealizations introduce approximation errors, they are not credible representations of behavior in all circumstances. In better textbooks, their use is typically restricted by a vague set of conditions and tacit experience. Consider the following from the standard reference for stress/strain equations [18, page 93], which is more precise than most texts:

> **7.1 Straight Beams (Common Case) Elastically Stressed**
>
> The formulas of this article are based on the following assumptions: (1) The beam is of homogeneous material that has the same modulus of elasticity in tension and compression. (2) The beam is straight or nearly so; if it is slightly curved, the curvature is in the plane of bending and the radius of curvature is at least 10 times the depth. (3) The cross section is uniform. (4) The beam has at least one longitudinal plane of symmetry. (5) All loads and reactions are perpendicular to the axis of the beam and lie in the same plane, which is a longitudinal plane of symmetry. (6) The beam is long in proportion to its depth, the span/depth ratio being 8 or more for metal beams of compact section, 15 or more for beams with relatively thin webs, and 24 or more for rectangular timber beams. (7) The beam is not disproportionately wide. (8) The maximum stress does not exceed the proportional limit.
>
> ...The limitations stated here with respect to straightness and proportions of the beam correspond to a maximum error in calculated results of about 5%.

Our goal in this research is to provide answers to the following questions:

1. *How are these conditions derived?* What is the process by which a model is converted to a simpler, idealized version? What are the principles behind the form and content of the standard textbook rules of thumb?
2. *What do these conditions mean?* For what "nearly straight", "disproportionately wide" beams will error begin to exceed 5%? How can the conditions be relaxed if only 20% accuracy is needed? What if 1% accuracy is needed?
3. *What is the best method by which an automated modeling system should determine when an approximate model is credible?* The answer to this may not necessarily be the same as the answer to question 1.

This paper examines these issues for algebraic and ordinary differential equation models of up to second order. It describes IDEAL, a system which generates idealized versions of a given model and provides measurable information about the model's error. The key enabler is recognizing the centrality of context in the idealization process - the idealizations that are generated and the limits that are placed on their use reflect the (intended) user's typical cases. We begin by describing how idealized models are derived. Section examines how approximation error should be managed in an automated modeling setting, while Section describes the principles behind the kinds of conditions stated above and a technique, called *credibility domain* *synthesis*, for generating them. It closes with a discussion of how the same functionality might be achieved for more complex systems. In particular, our ultimate goal is to be able to reproduce the above passage, which requires the analysis of 3-dimensional, 4th-order partial differential equations.

Idealizations

A *model* $\mathcal{M}$ contains a set of (algebraic and ordinary differential) equations E describing the behavior of some physical system in terms of variables $V = \{t, y_2, \ldots, y_{k-1}, p_k, \ldots, p_n\}$, where y_i represents a dependent variable, and p_i represents a constant, model parameter (i.e., p_i is a function of elements external to the model).[1] At most one varying independent variable t is allowed (which typically denotes time). Each model also has an associated set of logical preconditions, as described in [6]. We make the simplification that all analyses occur in a single operating region (i.e., the status of the preconditions does not change and thus can be ignored for the purposes of this paper).

A *behavior* is a vector $\mathbf{v} = [v_1, \ldots, v_n]$ of assignments to V as a function of t over the interval $t \in [0, t_f]$. A set of *boundary conditions* B specify values for t_f, the model parameters, and $y_i(0)$ for all y_i such that B and $\mathcal{M}$ uniquely specify a behavior $\mathbf{v} = \text{BEHAVIOR}(\mathcal{M}, B)$.

An idealized model $\mathcal{M}^*$ arises from the detection of order of magnitude relationships, such as those described in [10; 7], which enable the elimination of negligible terms. This can produce significant simplifications by reducing simultaneities and nonlinearities, and enabling closed-form, analytic solutions. In this paper, we consider the use of the two most common idealization assumptions:

$$\text{DOMINANCE-REDUCTION:} \quad A + B \approx A \quad \textit{given} \ |A| \gg |B|$$

$$\text{ISO-REDUCTION:} \quad \frac{dy}{dx} = 0 \quad \textit{given} \ \frac{\frac{dy}{dx}}{y} \approx 0$$

Dominance-reduction ignores negligible influences on a quantity and is the basis for idealizations like frictionless motion. When applied to derivative pairs, it offers one approach to time-scale approximation:

$$\text{TIME-SCALE-REDUCTION:} \quad \frac{dy_2}{dx} = 0 \quad \textit{given} \ |\frac{dy_1}{dx}| \gg |\frac{dy_2}{dx}|$$

Iso-reduction assumes constancy and is the basis for idealizations like quasi-statics, homogeneous materials, and orthogonal geometries. It is often the key enabler to obtaining analytic solutions.

In general, order of magnitude reasoning requires a carefully designed set of inference rules (e.g., approximate equality is not transitive [10]). For the class of ODEs currently being studied, algebraic operations across a set of equations are unnecessary and these issues do not arise. Thus, IDEAL currently uses only the two idealization rules without the associated machinery to propagate their consequences.[2]

Given that $\mathcal{M}^*$ is an idealization of $\mathcal{M}$, the error function $\mathbf{e}$ of $\mathcal{M}^*$'s approximate behavior $\mathbf{v}^*$ is measured with respect to $\mathcal{M}$'s predicted behavior $\mathbf{v}$ and couched in an appropriate scalar *norm* $\mathbf{e} = \| \mathbf{v}^* - \mathbf{v} \|$, the standard measure of goodness of an approximation [3]. The results are independent of the particular norm. In the examples we will use the maximum (L_∞) norm for the relative error

$$e_i(v_i) = \max_{t \in [0, t_f]} \left| \frac{v_i^*(t) - v_i(t)}{v_i(t)} \right|$$

where $\mathbf{e} = [e_1, \ldots, e_n]$ and e_i is the error norm for variable v_i. At times (section), the instantaneous error value as a function of t will be used in place of the absolute value norm.

A model is *credible* with respect to error tolerance $\tau = [\tau_a, \ldots, \tau_b]$ if $e_i(v_i) \le \tau_i$ for every variable v_i for which a tolerance has been specified. Not all variables need have an associated tolerance; the error of these variables is unmonitored and unconstrained. The *credibility domain* of an idealized model is the range of model parameters and t for which the simplified model is a credible representation of the original model [2].

Idealization process

A model may be idealized in at least two settings. In the on-demand setting, idealizations are made during the course of using a model to analyze a specific system. In the compilation setting, a model is idealized into a set of simpler models *a priori* by assuming the different potential relationships that enable use of the idealization rules. Because much of the on-demand setting is a special case of the compilation setting, we will focus solely on the latter. The key then is to pre-identify the set of enabling relationships that might arise. One straw-man approach would be to systematically explore all algebraic combinations and assume hypothetical situations in which the idealizations' conditions would hold. For example, for every pattern $A + B$, we could make one reduction based on $A \gg B$ and another based on $A \ll B$ (when consistent with the given equations). To make *useful* idealizations, we must have information about what relationships are possible or likely in practice. This is critical both in guiding the idealization process and in characterizing each idealized model's credibility domain (as discussed in section).

The more specific the information about what is likely, the more the idealized models may be tuned for one's specific future needs. Information about the population of analytic tasks is represented as a set of distributions over parameter values across problems and their variability within each problem (see Table 1).

[1] This is also known as an *exogenous* variable in the economics and AI literature. Throughout, we will try to use standard engineering terminology and indicate synonyms.

[2] The current implementation is in Mathematica, which is a hindrance to implementing the kinds of order of magnitude systems described in [10; 7].

Table 1: Some distributions characterizing an analyst's typical problem set.

Distribution of parameter values	
Simple Ranges	$p \in [0.1..1.5]$
Independent	uniform, normal, truncated
Joint (e.g., A may never be small when B is large)	
Distribution of function types	
Constant	$\frac{dy}{dx} = 0$
Nearly Constant	$\frac{dy}{dx} \approx 0$
Dependent	$y = y(x),\ \lvert\frac{dy}{dx}\rvert > 0$

Distributions on parameter values indicate which inequalities are likely or possible. They provide information about the population of tasks as a whole. Parameters specified by simple ranges are treated as having uniform distributions over their range. Distributions on function types provide information about the per-task behavior of parameters. For example, material densities may have wide variance across different analyses, but are normally constant throughout the material during any single analysis. These distributions are currently given as input; in the context of a CAD environment, they could easily be obtained by saving information about each analytic session.

IDEAL is guided by two factors - the form of the original equations and the problems to which they are typically applied. Given a model and associated distributions, it proceeds as follows:[3]

1. Syntactically identify candidates for reduction. Based on the two reduction rules, a candidate is either a sum or a derivative.
2. For sums, cluster the addenda into all possible dominant/negligible equivalence classes based on the given distributions. Each parameter's possible range is truncated at three standard deviations (otherwise it could be infinite and lead to spurious order of magnitude relationships).
3. For each possible reduction rule application, derive an idealized model under the assumption of the rule's applicability condition.
4. Repeat for consistent combinations of idealization assumptions (c.f., assumption combination in an ATMS [4]).

Example (sliding motion) Figure 1 illustrates the problem of determining the velocity and position of a block as it slides down an incline. The given model considers the influences of gravity, sliding friction, and air resistance. Due to the nonlinear response to air resistance, the model has no analytic solution.[4] The

[3] This is a more complete version of the algorithm described in [13]. For example, the earlier algorithm did not consider the ISO-REDUCTION rule.

[4] Well, at least not one that Mathematica can find.

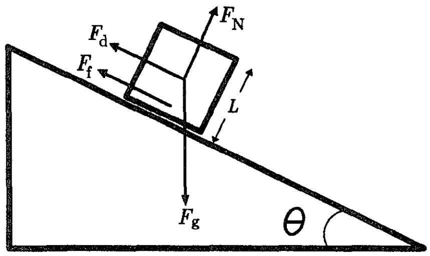

$\mathcal{M}_{gfd}$: $\frac{dv}{dt} = a_g + a_f + a_d$
$\frac{dx}{dt} = v$

Gravity: $a_g = g \sin\theta$
Sliding Friction: $a_f = -\mu_k g \cos\theta \mathrm{sgn}(v)$
Air Resistance: $a_d = -C_d \rho_{air} L^2 v^2 \mathrm{sgn}(v)/2M$

Distributions		
parameter	*type*	*pdf*
t	truncated normal ($t \in [0..\infty]$)	$\frac{1}{e^{\frac{t^2}{20} 4.52}}$
θ	uniform	$[30^\circ..60^\circ]$
μ_k	truncated, skewed ($\mu_k \in [0.2..0.55]$)	$\mu_k - \mu_k^3 + 2.54$
$\frac{dv}{dt}$	dependent	
$\frac{dx}{dt}$	dependent	

Figure 1: A block slides down an inclined plane. Need we model sliding friction, air drag, or both? In the table, pdf = *probability density function*

methods apply to higher-dimensions, but to enable 3-dimensional visualization and simplify the presentation, the initial velocity, $v_0 = 0$, and the air resistance coefficient ($C_d \rho_{air} L^2$) will be treated as constants.

IDEAL begins by identifying patterns for which the idealization rules may apply. In this case, there is the single sum

$$a_g + a_f + a_d$$

The assumption of $|A| \gg |B|$ is limited by requiring that at least $|A/B| \geq 10$ must be possible. Using this constraint and the given distributions, only one partial ordering is possible: $|a_g + a_f| \gg |a_d|$. This enables, via dominance-reduction, the derivation of a simple linear approximation $\mathcal{M}_{gf}$:

$$\frac{dv}{dt} = A_{gf} = a_g + a_f, \qquad \frac{dx}{dt} = v$$

from which we can derive

$$v(t) = A_{gf}\, t, \quad x(t) = \frac{A_{gf}}{2} t^2 + x_0 \qquad (\mathcal{M}_{gf})$$
$$\textit{assuming } A_{gf} \gg a_d$$

Had the distributions covered wider ranges for angle and time, and allowed air resistance to vary, a space of possible models, each with its own assumptions and credibility domain, would be derived. For example, high viscosity, long duration, and low friction would make the friction term insignificant with respect to the drag term, resulting in another idealized model:

$\frac{dv}{dt} = g \sin\theta - C_d \rho_{air} L^2 v^2 \mathrm{sgn}(v)/2M \qquad (\mathcal{M}_{gd})$
assuming $A_{gd} \gg a_f$

Error management for automated modeling

The idealized model $\mathcal{M}_{gf}$ derived in the example offers a considerable computational savings over its more detailed counterpart. Unfortunately, it is also quite non-operational as stated. What does $A_{gf} \gg a_d$ mean? When should one expect 5%, 10%, or 50% error from the model? What we would like is a mechanism for bounding the model's error that is (1) easy to compute at problem solving time – it should require much less time than the time savings gained by making the idealization, and (2) reliable – failure, and subsequent search for a more accurate model, should be the exception rather than the rule.

One appealing approach lay in the kinds of thresholds illustrated in the introduction, but augmented with some clarifying quantitative information. However, it is not as simple as deriving $\mathbf{e}$ as a function of A_{gf}/a_d or sampling different values for A_{gf}/a_d and computing the corresponding error. For a specified error threshold of 5%, the meaning of $A_{gf} \gg a_d$ is strongly influenced by the model parameters and the independent variable's interval. Figure 2 illustrates the error in position as a function of $A_{gf} = a_g + a_f$ and time t. The problem is further compounded in the context of differential equations. Not only does a_d change with time, the influence of error accumulation as time progresses can dominate that of the simple $A_{gf} \gg a_d$ relationship. Second, much of the requisite information cannot be obtained analytically (e.g., $\mathbf{e}(A_{gf}, t)$). For each value of A_{gf} and t_f, we must numerically integrate out to t_f. Thus, any mechanism for *a priori* bounding the model's error presupposes a solution to a difficult, N-dimensional error analysis problem.

The key lies in the following observation: only an approximate view of the error's behavior is needed – unlike the original approximation, this "meta-approximation" need not be very accurate. For example, a 5% error estimate that is potentially off by 20% means that the error may actually be only as much as 6%. This enables the use of the following simple procedure:

1. Sample the error's behavior over the specified distributions to obtain a set of datapoints.
2. Derive an approximate equation for each e_i as a function of the independent variable and model parameters by fitting a polynomial to the N-dimensional surface of datapoints.

If the error is moderately smooth, this will provide a very reliable estimate of the model's error.[5] For $\mathcal{M}_{gf}$'s

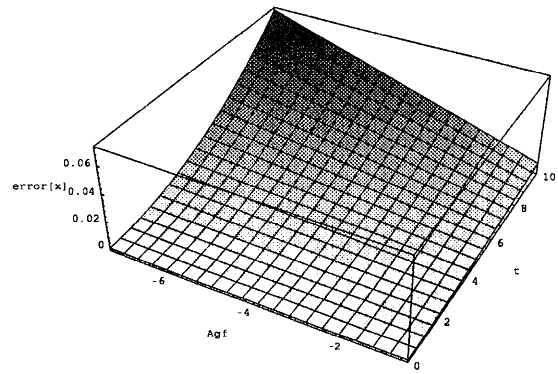

Figure 2: Percentage error in position x over the ranges for $A_{gf} = a_g + a_f$ and time t produced by $\mathcal{M}_{gf}$.

error in position x (shown in Figure 2), the resulting approximating polynomial is

$$\begin{aligned} e_x = {} & 3.52524\ 10^{-6} + 2.20497\ 10^{-6}\ A_{gf} - 3.18142\ 10^{-6}\ A_{gf}^2 \\ & -4.02174\ 10^{-7}\ A_{gf}^3 - 0.0000179397\ t \\ & -0.0000191978\ A_{gf}\ t - 1.20102\ 10^{-6}\ A_{gf}^2\ t \\ & +3.68432\ 10^{-6}\ t^2 - 0.0000940654\ A_{gf}\ t^2 \\ & +2.83115\ 10^{-8}\ A_{gf}^2\ t^2 - 1.1435\ 10^{-7}\ t^3 \end{aligned}$$

At this point, the specified requirements (easy to compute and reliable) have both been satisfied, without generating explicit thresholds! Although not as comprehensible, from an automated modeling perspective this approximate error equation is preferable because it provides two additional highly desirable features: (3) a continuous estimate of error that is better able to respond to differing accuracy requirements than a simple binary threshold, and (4) coverage of the entire problem distribution space by avoiding the rectangular discretization imposed by thresholds on individual dimensions.

Credibility domain synthesis

The question still remains – where do conditions like "the beam is not disproportionately wide" come from and what do they mean? They are clearly useful in providing intuitive, qualitative indications of a model's credibility domain. Further, for more complex systems, increased dimensionality may render the derivation of an explicit error function infeasible. The basic goal is to identify bounds in the independent variable and model parameters that specify a region within the model's credibility domain for a given error tolerance. This is the *credibility domain synthesis* problem:

find t_f *and* p_i^-, p_i^+ *for every* $p_i \in P$ *such that*
$$0 \le t \le t_f \ \wedge\ [\forall (p_i \in P), p_i^- \le p_i \le p_i^+] \ \rightarrow\ \mathbf{e} < \tau$$

[5] As one reviewer correctly noted, global polynomial approximations are sensitive to poles in the function being modeled. For the general case, a more reliable method is needed, such as local interpolation or regression on a more phenomena-specific basis function. There is nothing in the IDEAL algorithm which limits use of these methods.

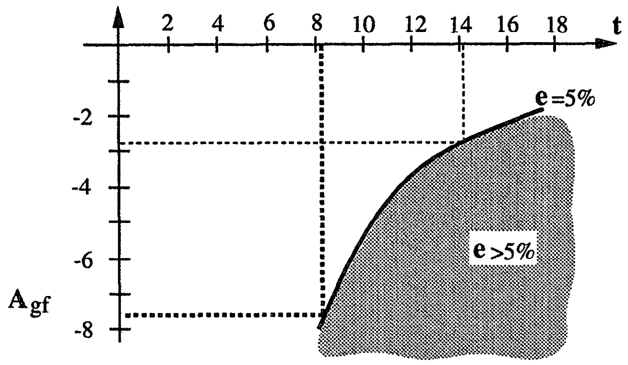

Figure 3: The error function imposes conservation laws on the shape of the credibility domain.

Unfortunately, these dimensions are interdependent. Increasing the allowable interval for p_i decreases the corresponding interval for p_j. Figure 3 illustrates for $\mathcal{M}_{gf}$ subject to a 5% error threshold. A credibility domain that maximizes the latitude for t also minimizes the latitude for A_{gf}. What criteria should be used to determine the shape of the hyperrectangle? Intuitively, the shape should be the one that maximizes the idealized model's expected future utility. We currently define future utility as its prior probability. Other influences on utility, when available, can be easily added to this definition. These include the cost of obtaining a value for p_i and the likely measurement error of p_i. Given distributions on parameter values and derivatives, the credibility domain synthesis problem can be precisely formulated as the following optimization problem:

minimize

$$F(t_f, p_k^-, p_k^+, \ldots, p_n^-, p_n^+) =$$
$$1 - P(0 \le t \le t_f,\ p_k^- \le p_k \le p_k^+, \ldots, p_n^- \le p_n \le p_n^+)$$

subject to $\mathrm{e} \le \tau$

For the case of $\mathcal{M}_{gf}$ and the distributions given in Figure 1, the optimal credibility domain is

$$t < 8.35 \ \wedge\ \mu_k > 0.2 \ \wedge\ \theta < 60°$$

which has a prior probability of 0.975.

This formulation has several beautiful properties:

1. The credibility domain is circumscribed by clear and easily computed conditions.
2. It maximizes the idealized model's future utility according to the user's typical needs.
3. It offers a precise explanation of the principles underlying the standard textbook rules of thumb.

In particular, it explains some very interesting aspects of the passage quoted in the introduction. For example, a careful examination of the theory of elasticity [14], from which the passage's corresponding formulas were derived, shows that several influences on the error are omitted. Why? How is that (likely to be) sound? Consider the conditions synthesized for $\mathcal{M}_{gf}$. The limits for μ and θ cover their entire distribution; they are irrelevant with respect to the anticipated analytic tasks and may be omitted.[6] Only the independent variable's threshold imposes a real limit with respect to its distribution in practice.

Like the textbook conditions, credibility domain synthesis makes one assumption about the error's behavior - it must not exceed τ inside the bounding region. This is guaranteed if it is unimodal and concave between thresholds, or, if convex, strictly decreasing from the threshold. $\mathcal{M}_{gf}$ satisfies the former condition. However, the current implementation lacks the ability to construct such proofs, beyond simply checking the derived datapoints.

Related Work

Our stance, starting with [13], has been that the traditional AI paradigm of search is both unnecessary and inappropriate for automated modeling because experienced engineers rarely search, typically selecting the appropriate model first. The key research question is then to identify the tacit knowledge such an engineer possesses. [5] explores the use of experiential knowledge at the level of individual cases. By observing over the course of time a model's credibility in different parts of its parameter space, *credibility extrapolation* can predict the model's credibility as it is applied to new problems. This paper adopts a generalization stance - computationally analyze the error's behavior *a priori* and then summarize it with an easy to compute mechanism for evaluating model credibility.

This is in contrast to much of the work on automated management of approximations. In the graph of models approach [1], the task is to find the model whose predictions are sufficiently close to a given observation. Search begins with the simplest model, moves to a new model when prediction fails to match observation, and is guided by rules stating each approximation's qualitative effect on the model's predicted behavior. Weld's domain-independent formulation [15] uses the same basic architecture. Weld's derivation and use of *bounding abstractions* [16] has the potential to reduce this search significantly and shows great promise. Like our work, it attempts to determine when an approximation produces sound inferences. One exception to the search paradigm is Nayak [9], who performs a post-analysis validation for a system of algebraic equations using a mix of the accurate and approximate models. While promising, it's soundness proof currently rests on overly-optimistic assumptions about the error's propagation through the system.

[6] This occurred because μ and θ were bound by fixed intervals, while t had the ∞ tail of a normal distribution, which offers little probabilistic gain beyond 2-3 standard deviations.

Credibility domain synthesis most closely resembles methods for tolerance synthesis (e.g., [8]), which also typically use an optimization formulation. There, the objective function maximizes the allowable design tolerances subject to the design performance constraints.

Intriguing questions

Credibility domain synthesis suggests a model of the principles behind the form and content of the standard textbook rules of thumb. Their abstract, qualitative conditions, while seemingly vague, provide useful, general guidelines by identifying the important landmarks. Their exact values may then be ascertained with respect to the individual's personal typical problem solving context. This "typical" set of problems can be characterized by distributions on a model's parameters which in turn can be used to automatically provide simplified models that are specialized to particular needs.

The least satisfying element is the rather brute-force way in which the error function is obtained. While it only takes a few seconds on the described examples, they are relatively simple examples (several permutations of the sliding block example described here and the more complex fluid flow / heat exchanger example described in [5]). The approach will likely be intractable for higher-dimensional systems over wider distributions, particularly the 3-dimensional PDE beam deflection problem. How else might the requisite information be acquired? What is needed to reduce sampling is a more qualitative picture of the error's behavior. This suggests a number of possible future directions. One approach would be to analyze the phase space of the system to identify critical points and obtain a qualitative picture of its asymptotic behavior, which can in turn suggest where to measure [11; 12; 17]. Alternatively, one could use qualitative envisioning techniques to map out the error's behavior. The uncertainty with that approach lies in the possibility of excessive ambiguity. For some systems, traditional functional approximation techniques might be used to represent the error's behavior.

Acknowledgments

Discussions with Colin Williams, both on the techniques and on programming Mathematica, were very valuable.

References

[1] Addanki, S, Cremonini, R, and Penberthy, J. S. Graphs of models. *Artificial Intelligence*, 51(1-3):145–177, October 1991.

[2] Brayton, R. K and Spruce, R. *Sensitivity and Optimization*. Elsevier, Amsterdam, 1980.

[3] Dahlquist, G, Bjorck, A, and Anderson, N. *Numerical Methods*. Prentice-Hall, Inc, New Jersey, 1974.

[4] de Kleer, J. An assumption-based TMS. *Artificial Intelligence*, 28(2), March 1986.

[5] Falkenhainer, B. Modeling without amnesia: Making experience-sanctioned approximations. In *The Sixth International Workshop on Qualitative Reasoning*, Edinburgh, August 1992.

[6] Falkenhainer, B and Forbus, K. D. Compositional modeling: Finding the right model for the job. *Artificial Intelligence*, 51(1-3):95–143, October 1991.

[7] Mavrovouniotis, M and Stephanopoulos, G. Reasoning with orders of magnitude and approximate relations. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, pages 626–630, Seattle, WA, July 1987. Morgan Kaufmann.

[8] Michael, W and Siddall, J. N. The optimization problem with optimal tolerance assignment and full acceptance. *Journal of Mechanical Design*, 103:842–848, October 1981.

[9] Nayak, P. P. Validating approximate equilibrium models. In *Proceedings of the AAAI-91 Workshop on Model-Based Reasoning*, Anaheim, CA, July 1991. AAAI Press.

[10] Raiman, O. Order of magnitude reasoning. *Artificial Intelligence*, 51(1-3):11–38, October 1991.

[11] Sacks, E. Automatic qualitative analysis of dynamic systems using piecewise linear approximations. *Artificial Intelligence*, 41(3):313–364, 1989/90.

[12] Sacks, E. Automatic analysis of one-parameter planar ordinary differential equations by intelligent numerical simulation. *Artificial Intelligence*, 48(1):27–56, February 1991.

[13] Shirley, M and Falkenhainer, B. Explicit reasoning about accuracy for approximating physical systems. In *Working Notes of the Automatic Generation of Approximations and Abstractions Workshop*, July 1990.

[14] Timoshenko, S. *Theory of Elasticity*. McGraw-Hill, New York, 1934.

[15] Weld, D. Approximation reformulations. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, Boston, MA, July 1990. AAAI Press.

[16] Weld, D. S. Reasoning about model accuracy. *Artificial Intelligence*, 56(2-3):255–300, August 1992.

[17] Yip, K. M.-K. Understanding complex dynamics by visual and symbolic reasoning. *Artificial Intelligence*, 51(1-3):179–221, October 1991.

[18] Young, W. C. *Roark's Formulas for Stress & Strain, Sixth Edition*. McGraw-Hill, New York, NY, 1989.