Modeling without amnesia: Making experience-sanctioned approximations

Brian Falkenhainer
Xerox Palo Alto Research Center
3333 Coyote Hill Road, Palo Alto CA 94304 USA

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. The central problem is determining when a model will be sufficiently accurate for a given task in a way that is simple and doesn't overwhelm the benefits of having a simplified model. This work presents *credibility extrapolation*, an inference procedure for using experience to sanction the use of approximations. It uses the accuracy measure obtained from prior (analytic or empirical) observations to project an accuracy bounds on the proposed model for a given setting *prior* to its use. This allows validation of the model without resorting to more expensive measures such as search or empirical confirmation. We then describe representation methods that make the storage and retrieval process efficient. The technique is illustrated on a moderate sized example.

1 Introduction

When using an approximate model to predict the behavior of some physical system, how can its approximations' appropriateness be determined? For some approximations, general-purpose rules of thumb may be consulted [1, 7] (e.g., Biot modulus, Mach number, etc. [10, 14]). Alternatively, a common answer to this question is that you should empirically or analytically confirm the derived result [3, 1, 17, 21, 13]. Though an important part of modeling, this isn't done for every behavior derived from the simplified model; there would be little point in having made the simplification. Some approaches suggest a search paradigm, trying one model and then another until a useful one is found [1, 22]. Yet, clearly an experienced engineer knows a lot about the available models and rarely searches. Missing from these accounts is this experience factor. One of an engineer's most oft used skills is the ability to retain and reason from past analyses and observations. Knowledge of a

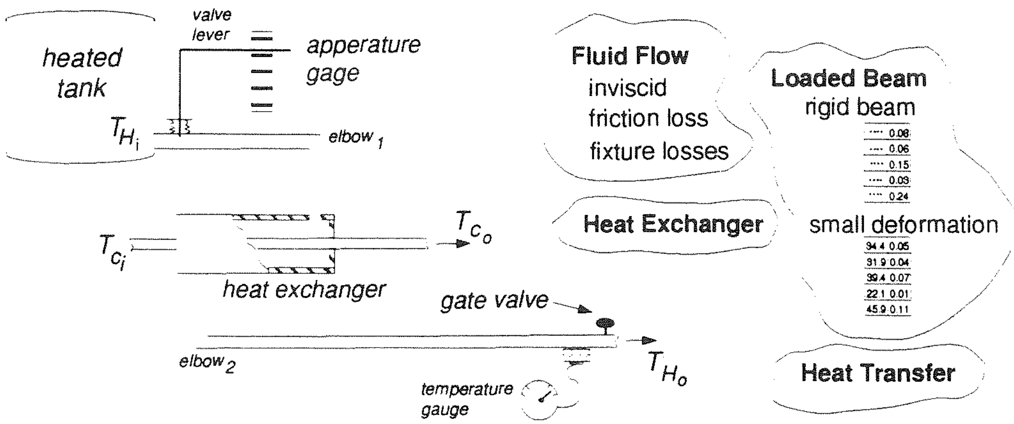

Figure 1: What is the temperature, within ±5%, at the gate valve outlet?

model's accuracy at some points can be used, often in quite simple ways, to bound its accuracy as it is applied to different parts of the behavior space.

This paper examines the issues that arise when trying to incorporate experience into an automated modeling framework. It describes methods for model representation that eliminate the complex indexing and retrieval problems normally associated with experiential systems. The core rule of inference, *credibility extrapolation*, uses the error measured (analytically or empirically) in prior situations to bound the error of the approximate model for the current problem description. In the case of a beam loaded at one end (the valve lever in Figure 1), if its deformation was found to be insignificant under a load of 5 N, then clearly, "all else being equal", it will be insignificant if the current problem specifies a load of 3 N and the beam may be safely treated as a rigid body. Further, this information applies to any situation in which an object is subject to the same distribution of forces and is not specific to the valve lever. No search or complex validation analysis is necessary. However, we must address the fact that rarely will all else be equal.

The paper begins by establishing some basic primitives and presenting the core credibility extrapolation inference procedure. Section 4 looks at representational and memory issues aimed at reducing complexity and increasing adaptability to new situations. The valved heat exchanger example is used for illustration.

2 Models and their accuracy

For what follows, it is important to understand the larger context. We assume an external automated modeling system designed to formulate an appropriate model to answer a given query (e.g., [7]). It must search a space of candidate models differing in perspective, resolution, and accuracy, weighing the relative costs and benefits in the process. This paper focuses solely on the accuracy dimension. The subtask being examined takes as given a relevant model and a query. Its function is to determine if the accuracy of the model satisfies the conditions imposed by the query.

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 [7]. 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 enough values for $y_i(t_j)$ such that B and $\mathcal{M}$ uniquely specify a behavior $\mathbf{v} = \text{BEHAVIOR}(\mathcal{M}, B)$.

In this paper, we treat errors as arising from only two types of approximation – physical *idealization* of the phenomenon (e.g., frictionless motion) and *mathematical approximation* of a function (e.g., linear discretization, polynomial expansion, etc).[2] An idealization assumes that some physical property is sufficiently close (in the current context) to a limit value (e.g., 0, 1, or ∞) to render the associated aspect of the model negligible and thus ignorable.[3]

The error function $\mathbf{e}$ of an approximate model $\mathcal{M}^*$ is taken to be with respect to some *base model* $\mathcal{M}^B$, which may denote the physical system being analyzed or an equational model we wish to approximate. A *base behavior* $\mathbf{v}^B$ will typically be referred to in a virtual sense as being that behavior which would be observed if measurements of $\mathcal{M}^B$ were taken. Drawing upon standard practice in numerical analysis for measuring an approximation's quality [3], we express the error function $\mathbf{e}$ of an approximate behavior $\mathbf{v}^*$ as an appropriate scalar *norm* $\mathbf{e} = \| \mathbf{v}^* - \mathbf{v}^B \|$. The results are independent of the particular norm. In the examples we will use the maximum (L_∞) norm:

$$e_i(v_i, \mathcal{M}^B) \;=\; \max_{t \in [0, t_f]} \mid v_i(x) - v_i^B(x) \mid$$

where $e_i \in \mathbf{e} = [e_1, \ldots, e_k]$ is the error norm for variable v_i.

Intuitively, a model $\mathcal{M}^*$ is an approximation of model $\mathcal{M}$ if $V^* \subseteq V$, E^* is simpler than E (e.g., of lower order), and critically, the error norm for $\mathcal{M}^*$'s behaviors will tend to be greater than for $\mathcal{M}$'s behaviors.

[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]Of course, errors can arise from many other sources as well, such as numeric instability and unexpected phenomena. We ignore these factors for the present.

[3]The idealization / approximation distinction is standard in the engineering literature. Idealizations with a specific property are called *fitting approximations* in [21], which shows how this property can be exploited to refine an approximate model found to deviate from observation.

A *query* $\mathcal{Q} = (W, Q, B, \mathcal{T})$ has four elements, where

- W is a set of variables. The model used to answer $\mathcal{Q}$ must satisfy $W \subseteq V$.
- Q is (1) a relation between variables in W (e.g., equality, inequality) whose status we wish to test or (2) a request for each variable in W as a function of t, for t in $[0, t_f]$.
- B is a set of boundary conditions.
- $\mathcal{T}$ is a vector of error tolerances $\mathcal{T} = [\tau_i, \ldots, \tau_k]$, where $\tau_j \in \Re^+$ specifies that $e_j(v_j, \mathcal{M}^B) < \tau_j$ must hold for every $v_j \in W$. This condition is abbreviated $\mathbf{e}(\mathbf{v}^*, \mathcal{M}^B) \leq \mathcal{T}$, where the tolerance for the unconstrained variables $V - W$ is taken to be infinity. The type and magnitude of the tolerance (e.g., absolute or relative) depends on the goals of the task. To simplify the remaining presentation, we will write $\text{VARS}(\mathcal{T})$ in place of W.

Tolerance $\mathcal{T}'$ is *weaker than* tolerance $\mathcal{T}$ if $\text{VARS}(\mathcal{T}') \subseteq \text{VARS}(\mathcal{T})$ and $\tau_i' \geq \tau_i$ for every $v_i \in \text{VARS}(\mathcal{T}')$. This relationship enables a more general formulation for comparing the results of different queries.

The standard definition of model credibility with respect to a tolerance and error norm [3] may be expressed as follows:

Definition 1 (model credibility) A behavior $\mathbf{v}^*$ is *credible* with respect to base $\mathcal{M}^B$ and tolerance $\mathcal{T}$, written $\text{CREDIBLE}(\mathbf{v}^*, \mathcal{M}^B, \mathcal{T})$, if $\mathbf{e}(\mathbf{v}^*, \mathcal{M}^B) \leq \mathcal{T}$. A model $\mathcal{M}$ is *credible* for query $\mathcal{Q} = (W, Q, B, \mathcal{T})$, written $\text{CREDIBLE}(\mathcal{M}, \mathcal{M}^B, \mathcal{Q})$, if $\mathbf{v}^* = \text{BEHAVIOR}(\mathcal{M}, B)$ and $\text{CREDIBLE}(\mathbf{v}^*, \mathcal{M}^B, \mathcal{T})$.

3 Credibility extrapolation

Our basic goal in evaluating the credibility of an approximate model for some query is to determine if $\mathbf{e} < \mathcal{T}$ holds. However, $\mathbf{e}$ can be quite expensive to compute. Luckily, this is often unnecessary. Given that a model's accuracy is known at some points, we can often use knowledge of the error's qualitative behavior to bound its accuracy at new points. We call this process *credibility extrapolation*. Importantly, only knowledge of the behavior of the error is required, not its specific valuation. In fact, it isn't necessary in principle that the error function itself be known (e.g., when there is no base equational model, only empirical data). The more that is known about its behavior, the greater the inferential power in determining the credibility of a model.

3.1 Proof by reduction

One approach is to directly compare the current prediction problem to a single prior behavior and show the credibility of $\mathcal{M}$ at $\mathbf{v}_{new}$ given the credibility of $\mathcal{M}$ at

$\mathbf{v}_{old}$. This approach is particularly suited to low-experience settings, as in when $\mathbf{e}$ is known at only a single point.

Proposition 1 (Reduction by measurement) *Let $\mathcal{T}$ be a tolerance vector. Let $\mathbf{v}'$ and $\mathbf{v}$ be behaviors derived from $\mathcal{M}$ such that the value of $e_i(v'_i, \mathcal{M}^B)$ is known for $v'_i \in$ VARS$(\mathcal{T})$. Let $\Delta\mathbf{e} = \mathbf{e}(\mathbf{v}, \mathcal{M}^B) - \mathbf{e}(\mathbf{v}', \mathcal{M}^B)$.*

If $\mathbf{e}(\mathbf{v}', \mathcal{M}^B) \leq \mathcal{T}$ and $\Delta\mathbf{e} \leq 0$, then CREDIBLE$(\mathbf{v}, \mathcal{M}^B, \mathcal{T})$.

This bases a model's credibility directly on the available datapoints and is the typical mode of inference. However, there may be situations in which only the relative error is known, such as when a model's credibility for behavior $\mathbf{v}'$ is asserted directly, without an explicit error value. Useful inferences may still be drawn using the following rule.

Proposition 2 (Reduction by transitivity) *Let $\mathcal{T}'$ and $\mathcal{T}$ be tolerances. Let $\mathbf{v}'$ and $\mathbf{v}$ be behaviors derived from $\mathcal{M}$. Let $\Delta\mathbf{e} = \mathbf{e}(\mathbf{v}, \mathcal{M}^B) \leq \mathcal{T} - \mathbf{e}(\mathbf{v}', \mathcal{M}^B) \leq \mathcal{T}$.*

If CREDIBLE$(\mathbf{v}', \mathcal{M}^B, \mathcal{T}')$, $\Delta\mathbf{e} \leq 0$, *and $\mathcal{T}$ is weaker than $\mathcal{T}'$, then* CREDIBLE$(\mathbf{v}, \mathcal{M}^B, \mathcal{T})$.

Note that by using reduction to a previous case, the credibility of $\mathcal{M}$ for query $\mathcal{Q}$ can be determined *prior* to consuming resources to derive a behavior using $\mathcal{M}$ simply by examining perturbations in $\mathbf{e}$ caused by changes in the boundary conditions. This is in contrast to even the use of standard applicability constraints (e.g., Biot modulus), which often can only be checked *after* spending resources to derive a behavior.

The task is then to determine the sign of $\Delta\mathbf{e}$ given some $\Delta\mathbf{v}$. This is a special-case of the problem of multivariate large-change sensitivity analysis [2, 19]. Note that the problem here is much simpler; we seek only the direction of change in $\mathbf{e}$, not a specific sensitivity coefficient.[4] There are a variety of methods, including qualitative comparative analysis [20], inequality reasoning over the interval $\Delta\mathbf{v}$ and the magnitudes of $\partial\mathbf{e}/\partial v_i$ (e.g., BOUNDER [16]), or examination of the system's higher-order derivatives (i.e., its Jacobian and Hessian matrices). Here we illustrate for the simple (but not uncommon) class of monotone error functions, in which $\mathbf{e}$ is monotone in each variable in $\Delta\mathbf{v}$.

Example The valve lever mechanism in Figure 1 can be modeled such that the valve aperture a is a function of the lever position p and deformation Δy

$$a = p + \Delta y$$

The deformation can be modeling in the following two ways. The first is a standard elastic model:

$$\Delta y = -\frac{kpL^3}{3EI}$$

[4] For many cases, this computation is quite trivial. However, it can become quite complex. Being able to predict its efficiency is an important, but open problem.

The second ignores its deformation under load and makes a rigid body idealization:

$$\Delta y = 0$$

The measured error in the estimation of valve aperture using the rigid body idealization is 5% for $p = 0.15$. Suppose we wish to analyze the same situation with a new position $p = 0.12$. In this simple case, a purely qualitative argument suffices:

$$e \propto_{Q^+} p \ \wedge \ \Delta p < 0 \ \rightarrow \ \Delta \mathbf{e} < 0$$

Because $\mathcal{M}$ was credible in the first analysis and $\Delta \mathbf{e} < 0$, $\mathcal{M}$ is credible for the current analysis as well.

4 Model Composition and Reuse

Credibility extrapolation's efficiency and ability to flexibly apply experience in one setting to different settings is largely affected by the models' structure. One enhancement is to split up the problem and convert an analysis of one large model into separate analyses of its pieces. We do this by placing it in the context of *compositional modeling* [6, 7], a framework for organizing and formulating appropriate models in response to analytic goals. Another way is to reduce the model's dimensionality (the size of V) by precomputing algebraic aggregates $V' = f(V)$, $|V'| < |V|$ and performing sensitivity comparisons using the more compact model $F'(V')$. These topics are discussed in the following subsections.

4.1 Compositional models

A model and the physical system it models is typically an interconnection of primitive *modules* $C = \{c_1, \ldots, c_n\}$, where a module may be a component, process, bond-graph element, or other conceptual primitive phenomenon. Explicit recognition of this in the modeling framework is important for flexibility, reuse, and storage efficiency. Though not further decomposable, physically a module may correspond to some aggregate or arrangement of parts. We compose a model $\mathcal{M}$ from a set of *model fragments* $M = \{m_1, \ldots, m_n\}$, each of which contains a set of equations (partially) describing some module. For example, one model fragment may represent the general Bernoulli equation for flow through some plumbing, while others may specify different functions for the frictional losses within that flow. We require that each model fragment be associated with a single module such that there exists an onto function $F : M \rightarrow C$. In the valved heat exchanger example, the modules are the tank, levered valve, plumbing$_1$, heat exchanger, plumbing$_2$, and temperature gauge (see also Figure 2).

Model fragments define the granularity of the modeling building blocks. This decomposition provides significant leverage in two ways. First, it decomposes experience with one large model into experience with its constituent model fragments,

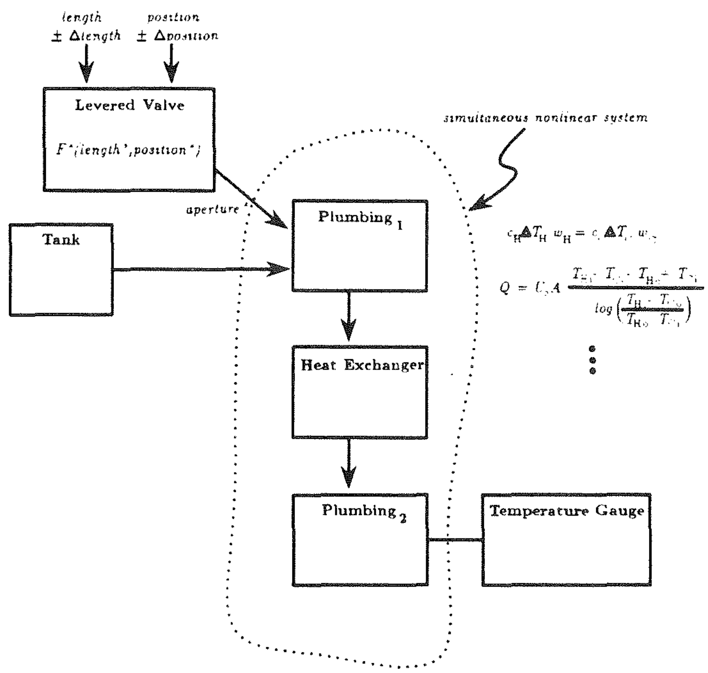

Figure 2: Error in a module's output results from transmission of input error and generation of internal error.

each of which may be subsequently used in contexts very different from the setting in which the experience was acquired. For example, information gained about the error function for the rigid-body idealization of the levered valve is independent of the specific heat exchanger context in which the valve is currently placed (although its effects on other parts of the system are of course dependent on context). Second, it decomposes error analysis; the composite error of a model $\mathcal{M}$ is a function of the internal errors introduced within each module's model and the external errors transmitted through each module's model. For example, in Figure 2, the levered valve module passes input errors (e.g., the length of the lever) on to its outputs. This error propagation is performed using standard error analysis techniques [3, 18]. At the same time, its idealized model F^* introduces new, additional errors to its outputs. This error is bounded using credibility extrapolation.

4.2 Case storage and retrieval

Structuring the domain theory in this manner makes the issues surrounding the storage and retrieval of experience straightforward. Importantly, much of the requisite storage, retrieval, and adaptation functionality comes at low computational cost.

Storage of a validated analysis (i.e., one whose error has been explicitly evaluated), is simple. First, the scenario is maximally decomposed into a set of linearizable aggregate modules. Typically, this decomposition will correspond to the scenario's original decomposition. However, for some scenarios, there may be modules' which are not separable (e.g., they form a system of nonlinear simultaneous equations). For the valved heat exchanger, the decomposition results in the following modules:

$\{tank\}$ $\{lever\ valve\}$ $\{plumbing_1,\ heat\ exchanger,\ plumbing_2\}$ $\{temperature\ gauge\}$

Because of the interdependencies between the mass flow and heat transfer rates of the plumbing and heat exchanger, these modules cannot be separated.[5]

Second, the behavior (and associated error) of each aggregate module is stored with the module definition and the conjunction of the modeling assumptions associated with that module's model fragments.

Case retrieval has two aspects. First, given a query and a candidate model for answering it, find a set of relevant cases from which to extrapolate the model's credibility. This too is simple because prior usage is stored directly with the modules. Second, from the set of relevant cases, select one from which to bound the model's error. For this we draw upon the intuition that from the current (unknown) point on the error surface, we want to look "uphill" ($\partial \mathbf{e}/\partial \mathbf{v} > 0$) to find cases that might have greater error. We seek the point $\mathbf{v}'$ that minimizes the weighted Euclidian distance

$$\sqrt{\sum_{v_i \in \mathbf{V}} \mid v_i' - v_i \mid^2 w_i}$$

where $w_i = 1$ if $\partial \mathbf{e}/\partial v_i \geq 0$ and $w_i = 0.5$ if $\partial \mathbf{e}/\partial v_i < 0$.[6] More experimentation is needed to evaluate the impact of specific weight assignments. If the error function is known to be monotone in each of its variables, then an even simpler technique is possible. For each point $(d_1, \ldots, d_n)$ at which the error is known, produce a rule of the form

$$\textit{If}\ \ v_1\ op_1\ d_1\ \wedge \ldots \wedge\ v_n\ op_n\ d_n \quad \rightarrow \quad \mathbf{e}_{new} \leq \mathbf{e}_{old}$$

where op_i is $\leq$ ($\geq$) if $\mathbf{e}$ is monotone increasing (decreasing) in v_i.[7]

4.3 Changes in topology

Credibility extrapolation as presented in Section 3 assumed a fixed set of equations; only the model parameters are allowed to change. While useful, it captures a somewhat limited view of experience. To the degree possible, we would like to extend experience with one situation to distinct, yet similar situations.

This is done in two ways. First, as described above, experience is tied to the modules and their model fragments, not the complete scenario model. Second, we

[5]A standard Newton-Raphson iterative technique is used to compute these values, while the Jacobian of the system is used to propagate input error through the system.

[6]This selection of weights is arbitrary, but matches the intuition that more weight should be given to dimensions in which the error is known to decrease. A more principled weighting is being sought.

[7]Currently unimplemented.

$$\mathcal{M}_{gfd} \qquad\qquad \mathcal{M}_{gf}$$

$$\frac{dv}{dt} = a_g + a_f + a_d \qquad \frac{dv}{dt} = a_g + a_f$$

$$\frac{dp}{dt} = v \qquad\qquad \frac{dp}{dt} = v$$

Gravity: $a_g = g \sin\theta$

Sliding Friction: $a_f = -\mu_k g \cos\theta \mathrm{sgn}(v)$

Air Resistance: $a_d = -C_d \rho_{air} L^2 v^2 \mathrm{sgn}(v)/2M$

Prior Data

	t_f	p	v	a
$\mathcal{M}_{gfd}$	6	15.2074	-4.9167	-0.81007
$\mathcal{M}_{gf}$	6	15.1648	-4.9451	-0.82418
e		-0.0426	0.0284	0.01411

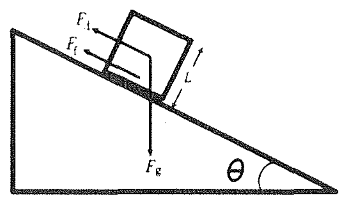

Figure 3: A block slides down an inclined plane. Need we model sliding friction, air drag, or both?

use *influences* to specify elements of *composable functions* (e.g., $\sum$, $\prod$) in a modular manner [9, 7]). For example, the frictional (head) loss in the plumbing due to elbows and other fixtures was modeled as a summation of the individual loss introduced by each fixture. Influences provide modularity at the function level; the effects of including or ignoring an influence can be determined in cases where no experience with that specific influence exists and enables credibility extrapolation even when the scenario topology changes. For example, experience obtained during analysis of a system containing four elbows can be used to place a bound on the error due to ignoring fixture losses in the two elbow configuration shown in Figure 1. Because these frictional losses are combined via influences, and there are fewer in this scenario, we can determine that the error due to ignoring losses in these fixtures will be less (given a lower or equal flow velocity), even though the configuration of fixtures is different.

4.4 Feature reduction

As a second example, consider the case of a block sliding down an incline (Figure 3). If drag was found to be insignificant at 20 km/h, then clearly, "all else being equal", it will be insignificant if the current problem specifies a speed of 10 km/h. This problem can be posed as the task of determining $\Delta\mathbf{e}$ given changes to the three variables $[a_g, a_f, a_d]$. However, it could just as easily be posed as the task of determining $\Delta\mathbf{e}$ given changes to the ten variables $[W, L, H, \rho_{block}, \theta, g, \mu_s, C_D, \rho_{air}, v]$. Each time a comparison to experience is made, we must compare a vector $\Delta\mathbf{v}$ of length N. Much depends on N. First, the complexity of sensitivity analysis in credibility extrapolation goes up with N. Second, the likelihood of a definitive credibility evaluation diminishes with N, due to increased chance of unresolved ambiguity. Finally, if approximate credibility evaluation algorithms are used (e.g., curve fit), the likelihood of an incorrect result goes up with N. Thus, it is important to minimize N. We do this reduction via algebraic aggregation of the variables into fewer terms.

For example, the block's acceleration (Figure 3) can be computed in several ways:

$$\begin{aligned}
a &= f_1(W, L, H, \rho_{block}, \theta, g, \mu_s, C_D, \rho_{air}, v) \\
a &= f_2(A, M, \theta, g, \mu_s, C_D, \rho_{air}, v) \quad (A = WH, M = LWH\rho_{block}) \\
a &= f_3(a_g, a_f, a_d)
\end{aligned}$$

Using f_1, the change in error due to an increase in block density and a decrease in block length would have to be resolved. Alternatively, using f_2, a change in density and length results in a single (potentially zero) change in mass M – a much easier problem. In general, minimizing N is a significant step in addressing the "all else being equal" problem mentioned in the introduction and, in the same vein, is of course one of the main points of dimensional analysis and the theory of similitude [11, 23].

Finding the optimal reformulation of the variables is beyond the scope of this paper. However, two simple reduction rules have provided significant leverage. First, if the function is a summation of influences, we treat each influence as a separate parameter to the function (c.f., f_3). This is why example 1 was posed as a change to the three variables $[a_g, a_f, a_d]$. Second, we aggregate model parameter products for which the function is invariant (i.e., $f(p_1, p_2) = f(p_2, p_1)$ or $f(p_1, 1/p_2) = f(1/p_2, p_1)$). Thus,

$$f(M, E, I, dv/dx) = M[1 + (dv/dx)^2]^{3/2}/EI$$

would become

$$f'(S, dv/dx) = S[1 + (dv/dx)^2]^{3/2}$$

where $S = M/EI$.

5 Related Work

This work builds directly from prior work with Shirley [17] and Forbus [7] and is intended to complement some of the existing approaches to reasoning about approximate models. In discrepancy-driven refinement [1, 17], an observation is used to find a model that fits. Credibility extrapolation can then tell if the model can still be used in subsequent analyses, thus allowing it to be used for prediction. Similarly, one could use bounding abstractions [22] to identify models that overestimate and credibility extrapolation to provide a bounds on the amount of overestimation and thus reduce or eliminate search.

Our approach is much like the caching that occurs in an ATMS [4], in which partial results obtained in one context are transparently reused in other contexts due to the underlying granularity of assumption maintenance. The complex memory and mapping issues that dominate treatments of analogy and case-based reasoning [8, 12, 15] do not arise. We achieve indexing and adaptation via the granularity inherent in the domain theory (the model fragments), rather than via partial situation matching and similarity-driven adaptation procedures.

6 Discussion

Analysis grounded in experience is fundamental to modeling due to the inherent uncertainty of equational models and the numerical methods which manipulate them. Given prior use of an approximate model, credibility extrapolation provides a mechanism for bounding the error of a model prior to its use. Further, by explicitly matching on model fragments, it eliminates the complex indexing and retrieval normally associated with experiential reasoning approaches like analogy and case-based reasoning. At the same time, the model fragments' granularity enables transfer of experience to different settings that make use of the same model fragments.

Credibility extrapolation allows for the observation to be generated analytically by sampling points in the more accurate model's domain (presumably at extrema conditions). A companion technique, called *credibility domain synthesis* [5], takes this to its natural extreme. It generates approximations of a given model by analyzing its domain and identifies explicit constraints for the most prevalent situations based on variance. Current work includes examining their relative merits.

The approach is currently implemented in pieces, with the model representation and storage occurring in our existing implementation of compositional modeling and most of the analysis occurring in Mathematica. I am currently working to connect the two systems.

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. A look at idealization. In *AAAI-92 workshop on approximation and abstraction of computational theories*, July 1992.

[6] Falkenhainer, B and Forbus, K. D. Setting up large-scale qualitative models. In *Proceedings of the Seventh National Conference on Artificial Intelligence*, pages 301–306, St. Paul, MN, August 1988. Morgan Kaufmann.

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

[8] Falkenhainer, B, Forbus, K. D, and Gentner, D. The structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, 41(1):1–63, November 1989.

[9] Forbus, K. D. Qualitative process theory. *Artificial Intelligence*, 24:85–168, 1984.

[10] Fox, R. W and McDonald, A. T. *Introduction to Fluid Mechanics.* John Wiley & Sons, New York, NY, third edition, 1985.

[11] Huntley, H. E. *Dimensional Analysis.* MacDonald, London, 1952.

[12] Kolodner, J (Ed.). *Proceedings of the DARPA Case-Based Reasoning Workshop.* Morgan Kaufmann, San Mateo, CA, 1988.

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

[14] Popov, E. P. *Mechanics of Materials.* Prentice-Hall, Inc, New Jersey, 1976.

[15] Riesbeck, C and Schank, R. *Inside Case-Based Reasoning.* Lawrence Erlbaum Associates, Hillsdale, NJ, 1989.

[16] Sacks, E. Hierarchical reasoning about inequalities. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, pages 649–654, Seattle, WA, July 1987. Morgan Kaufmann.

[17] 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.

[18] Taylor, J. R. *An Introduction to Error Analysis.* Oxford University Press, Mill Valley, CA, 1982.

[19] Tomovic, R and Vokobratovic, M. *General Sensitivity Theory.* Elsevier, New York, 1972.

[20] Weld, D. Comparative analysis. *Artificial Intelligence*, 36:333–373, 1988.

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

[22] Weld, D. S. Reasoning about model accuracy. Technical Report 91-05-02, Department of Computer Science, University of Washington, June 1991.

[23] Welty, J. R, Wicks, C. E, and Wilson, R. E. *Fundamentals of Momentum, Heat, and Mass Transfer (second edition).* John Wiley & Sons, New York, NY, 1984.