Meta Models to Aid Planning of Intelligent Machines

Paul K. Davis, senior scientist, RAND and professor, RAND Graduate School
and James H. Bigelow, senior mathematician, RAND
1700 Main St., Santa Monica, CA 90407-2138
pdavis@rand.org; James_Bigelow@rand.org

Abstract—Relatively simple low-resolution models are needed by human planners and probably by intelligent machines. Ideally, these should be high-level models developed in a multiresolution, multiperspective modeling (MRMPM) framework. That, however, is often difficult. We ask whether statistical meta modeling (i.e., development of response surfaces) can provide good low-resolution models if one already has a credible higher-resolution base model. We ask how meta models compare if they are derived from pure statistical methods, from a phenomenology-rich theoretical approach, or from some synthesis. To sharpen issues and generate insights, we have worked through a particular problem in detail. Our conclusions are generally negative about "purist" statistical meta models, which have serious shortcomings in explanatory power, in variance, and in ability to predict and explain the relative importance of contributing variables. Purely theoretical approaches, however, are often very difficult and not transparent. Fortunately, a synthesis of methods is feasible and likely to be fruitful. Some tentative principles are that: (1) a thoughtful "first-order" theoretical analysis conducted with MRMPM principles in mind can identify "aggregation fragments" to be used as variables in generalized regression and (2) this can also suggest structures to impose on the meta model that will assure dependences known to be important. Imposing such a structure can, e.g., assure that a meta model will predict failure of a system if any of its critical components fail. The theory-enhanced statistical meta model may also be much better than a naive statistical meta model in representing a system's performance when a competitor is systematically looking for a circumstances that will defeat the system. In that case, variables that are mathematically independent may be said to be strategically correlated. Although tentative, the suggested principles appear consistent with experience in theoretical and experimental physical science.

Index Terms— Multiresolution modeling, variable resolution modeling, response surfaces, meta models, model abstraction, planning models.

I. INTRODUCTION

This paper addresses the problem of how to develop low-resolution, meta models as part of a multiresolution family. In particular, it compares approaches based on phenomenological modeling with methods based on statistical methods. It then suggests some steps toward synthesis.

The paper begins with some background on multiresoluiton modeling and the reasons meta models are needed. It then discusses the ideal for phenomenological multiresolution modeling, which involves pure hierarchies. Although that ideal can sometimes be realized with considerable payoff, reality is often much more complex. As a result, developing phenomenology-driven multiresolution families proves quite difficult. This causes us to be interested in shortcuts, such as using statistical methods to develop meta models. The remainder of the paper is about our efforts to think about how statistical methods and more phenomenology-rich methods relate to each other and whether there is the possibility of combining features of both. We describe our initial hypotheses on the matter, the research approach we have taken so far, and observations to date.

II. BACKGROUND

A. Planner Needs for Low Resolution Models

It is well recognized by now that intelligent systems need planning modes in which they are able to recognize and compare alternative courses of action.[1-4] This planning requires a *broad* form of testing—i.e., the courses of action need to be evaluated for a wide range of circumstances. This is the domain of *exploratory analysis*, rather than the domain of refinement. The objective is often the classic goal of satisficing—finding a course of action that will "do the job," not necessarily optimally, but well enough.

It follows that humans, at least, typically need low-resolution models for planning. This is not simply a matter of saving time or money, but rather due to the human need to *understand* the basis for choosing one course of action over another, and to communicate that rationale to others—perhaps to persuade, or perhaps to convey a clear sense of mission intent. This need might not exist if a perfect model existed with perfect data, and if everyone accepted whatever the model said. That situation, however, rarely arises in higher level planning.

A corollary is that the need for simple, low-resolution models will continue to exist regardless of increasing computer speed. The need is fundamental. It is tied to the limits of cognition and curse of dimensionality.

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE **SEP 2001**	2. REPORT TYPE	3. DATES COVERED **00-00-2001 to 00-00-2001**

4. TITLE AND SUBTITLE **Meta Models to Aid Planning of Intelligent Machines**	5a. CONTRACT NUMBER
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **RAND Graduate School,1700 Main St,Santa Monica,CA,90407-2138**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
Proceedings of the 2001 Performance Metrics for Intelligent Systems Workshop (PerMIS ?01), Mexico City, Mexico on 4 Sep 2001

14. ABSTRACT
see report

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**8**	

It might be speculated that intelligent machines can be different on such matters. They have no emotional need for explanation and they may not need to explain their reasoning in simple terms—at least when communicating with other intelligent machines. Nonetheless, it seems likely that when the intelligent machines have imperfect models, limited data, and uncertainty about prospective operating conditions, they will suffer the same problems of bounded rationality addressed famously by the late Herb Simon[5] a half century ago. If so, they will also need simple, low-resolution models.

This said, even those who gravitate toward simple, low-resolution models will agree that to be useful, such models need to be grounded in reality. It is frequently easy to concoct plausible and attractive simple models, but such models are often flawed—so much so as to be counter productive. Sound "simple" models should be rooted in higher-resolution work. Thus, to conclude that the planning function requires simple models leads in due course to the requirement for multiresolution modeling (MRM). Indeed, it is not just a matter of resolution. Substantially different representations of reality (different "perspectives") may be essential in order to understand different facets of the underlying phenomenon or to make effective use of diverse forms or empirical data. Thus, what is needed is actually multiresolution, multiperspective modeling (MRMPM).

For the remainder of this paper we shall focus on MRM, but the more encompassing concept of MRMPM is important to keep in mind.

Having established motivation, let us now discuss what is involved in MRM.

B. Idealized Multiresolution Modeling: the Role of Hierarchies

For a phenomenologist, at least, the natural way to proceed in developing an MRM family is to design hierarchically.[6,7] Figure 1 illustrates schematically an idealized construct. One has only a few top-level variables (those in the low-resolution model), but each of these is determined by higher-resolution phenomena. The next level of detail will be a model with more variables and it, in turn, will depend on events at still higher detail. In Figure 1, the resulting hierarchical trees are pristine.

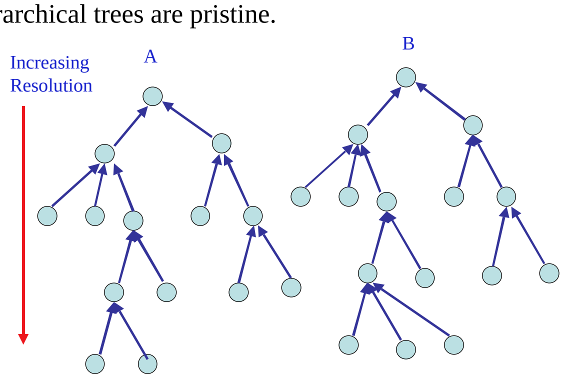

Figure 1—Idealized Multiresolution Modeling

Why is this "ideal?," or at least very desirable? For one thing, given such a multiresolution family, one can start at the top and then—as necessary—zoom to a higher level of detail, perhaps on only one part of the problem. For example, one might thoroughly understand variable A, but variable B might be uncomfortably abstract. If so, one could go down one or more levels of detail until the variables used are comfortable and sufficient—perhaps because they are explicitly tied to familiar empirical information. This zooming, however, would be on an as-necessary basis. Reasoning could be accomplished at as high a level, and with as few variables, as needed for comfort.

Such a multiresolution family would relate the microscopic and macroscopic worlds. It would provide a strong sense of "understanding" and the capacity to use diverse types of information. This relating of levels would not just be a matter of hand-waving. Instead, Figure 1 suggests that to establish good values for the higher-level variables when they are used as independent variables (inputs), one should conduct systematic experiments exercising the next higher-resolution model to generate appropriate "averages." Such experiments should be conducted over the entire n-dimensional space spanned by the independent variables of the higher resolution model. In some contexts, that is appropriately called a "scenario space."

Interestingly, the result of such calibration should generally be to produce *stochastic* variables. That is, if the higher-level (lower-resolution) model has two variables X and Y, and if we want to establish what reasonable values of X and Y might be, we should ordinarily expect that X and Y will need to be stochastic because of hidden variables.

Such idealized modeling is possible in many cases—if one thinks about doing it. Figure 2 shows an example drawn from recent defense work. [8]It shows the design of a module dealing with command and control issues in the evaluation of long-range precision fires. This model allows users to input directly the impact time of a weapon (measured relative to the ideal time of arrival at a target). This is often a useful quantity to parameterize and vary. However, the model also allows the user to work with more detailed variables as inputs. The second level of detail involves the descent time of the weapon (the time between when the weapon does its final target acquisition and tracking, when it is overhead, and when the weapon impacts) and the standard time-of-arrival error measuring the variation due to imperfect guidance system. At the most detailed level, the user must input the weapon's flight time, the delay between the receipt of sensor data on targets and the time that the data was valid, and so on.

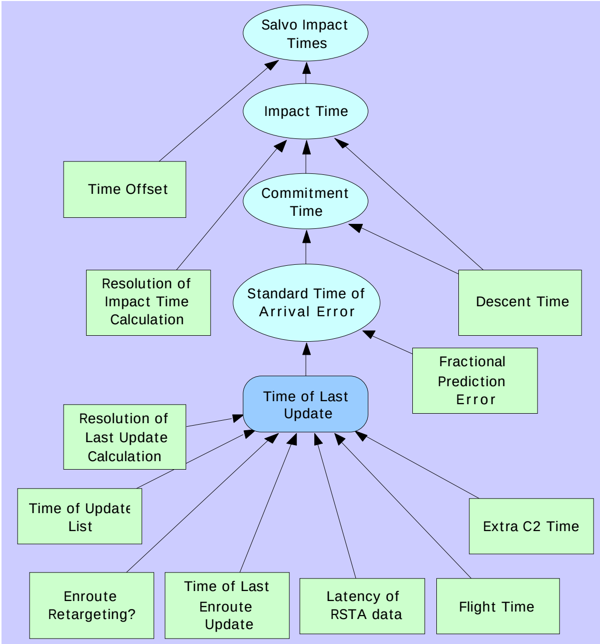

Figure 2—An Example of MRM Design

Idealized hierarchical design is unusual. If we look at an existing model and depict its relationships graphically, a more typical picture would be as in Figure 3. Here we see a good deal of cross talk and breakdown of the hierarchies. A common observation here is "Everything is connected to everything." Often, it is not evident how to simplify to something more like Figure 2.

This may be puzzling to those who know about and accept the principle that natural complex adaptive systems typically manifest the principle of nearly decomposable hierarchy:[5] that is, when viewed in the right way, the system can be decomposed into modules that interact only weakly. Such a decomposition is typically not evident when viewing the structure of existing complex models. Nor is it evident in freshly built models designed bottom-up with the common ethic of achieving verisimilitude. Indeed, it is not evident even in models built top-down if the designer is taking pains to include interactions that appear important. There are at least two points here. First, people only seldom design models with an image such as Figure 1 as a goal. Second, even if they try, they will find that their diagrams become muddled, as in Figure 3.

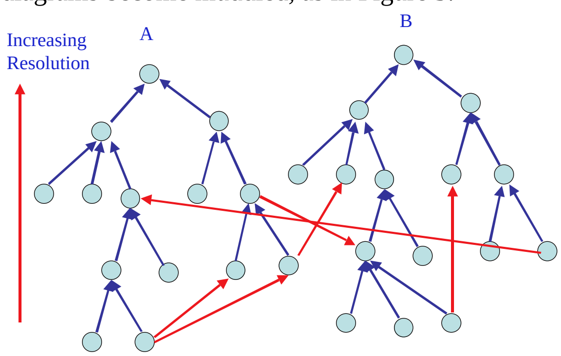

Figure 3—A More Typical Model Schematic

The solution, it might seem, is to recognize that approximations can eliminate the ugly interactions. Indeed, if one is willing to introduce approximations, then it *is* often possible to move much closer to the MRM ideal. And, if one does this right, one will rediscover the principle of nearly decomposable hierarchy.

C. *Intrusion of Reality*

Unfortunately, another fundamental reality intrudes here. The critical approximations are often valid only in limited domains. As one moves from one domain to another, the appropriate approximation may change drastically—not just through a change in some constant, but in the analytical structure. For example, aerodynamic drag may vary in one regime in proportion to an object's speed, whereas in another regime it may vary inversely with that speed. Yes, approximations are essential, but we should not expect to find simple, stable, universal approximations.[7]

The significance of this is that—once again—anyone attempting to develop a phenomenology-based MRM design in a given problem should not be surprised to find difficulties—difficulties great enough to comprise a PhD dissertation.

How, then, do we humans "get along" in this complex world? In fact, we do reasonably well. However, we are constantly changing the frames in which we operate (the approximate depictions of the world that allow us to reason and act). We do this so seamlessly that we often are not even aware that we have changed frames. The attribute of being able to carry along contradictory ideas at the same time—most celebrated in discussion of eastern philosophy, but actually a universal attribute—is arguably a manifestation of this.

What about machines? How will intelligent machines develop the diverse frames and skills to adopt the right frame at the right time? This remains very much a research question.

To complete our background discussion, let us summarize by observing that while simple, low-resolution models are needed, and while they need to be rooted in a multiresolution framework, achieving one is often difficult. Learning how to achieve MRM structures efficiently would be very desirable.

III. Can Statistical Meta Modeling Provide a Shortcut?

A. *General Issues*

The difficulties to which we have alluded so far are all tied to attempts to build phenomenological models—i.e., models rooted in theory and attempting to describe causes, effects, and other relationships. Suppose, however, we back away from this and ask whether an alternative approach is possible. The most obvious is statistical meta modeling, the very purpose of which is to develop simple "models" that represent well the behavior of systems on which some kind of data exists. The system in question may be a physical system and the data may be empirical. Alternatively, the system may be a detailed model (e.g., a simulation of a system) and the "data" may be outcomes of simulation runs. In some instances, the detailed models are large, complex, impenetrable, fragile, and slow. In other cases, they may be virtuous in all respects other than requiring expensive care and feeding. Typically, the base models are imperfect, with both known limits of applicability and errors.

In all of these cases, one can apply well known statistical methods to generate meta models. If a reasonably well accepted detailed model exists, why should we not adopt these methods to generate the simple, low-resolution models needed for planning?

This is the question we have been studying. We have sought to understand better the strengths and weaknesses of the phenomenological approach and the approach of statistical meta modeling. And we have sought opportunities for synthesis.

B. *An Aside*

One reason that pursuing this matter was of interest is that it highlights a substantial cultural divide, which can be characterized—with literary license—as follows. Suppose we ask whether using statistical methods to generate simple low-resolution models for planning is sensible. The responses from Cultures A and B might be::

> *Culture A*: "Of course they make sense; all that matters is representing behavior of the base model. I don't even *want* to understand the black box." (statisticians, some operations researchers, many social scientists,…?)
>
> *Culture B*: "No no no; the simple model should be a <u>model</u>, not some lousy regression. I'd rather calibrate a model that makes sense than work with a mysterious black box." (physical scientists, engineers,…?)

Culture A and Culture B even mean quite different things by the word "model." Fortunately, translations are possible.

IV. Approach

In our first assault on the issue, we proceeded on two tracks. On the first track, we theorized in the abstract, using simple examples to help, but without attempting anything rigorous. The purpose was to generate hypotheses for experiments. For our second, experimental, track, we decided to work though a particular nontrivial example drawing on a currently interesting military problem with which we were familiar. For that second track, we decided to

1. Construct (by embellishing an existing model) a complex, nonlinear model that we would treat as correct
2. Use standard methods to develop statistical meta models
3. Throw different degrees and types of theory at the problem—providing "hints" before applying the statistical apparatus.
4. Observe, compare results with differing levels of theory, compare results with expectations from initial notions, and learn.

More ambitious theoretical work would certainly be possible, but this hands-on experimentation was suitable to our state of knowledge and the limited time available for the research (in between our principal research efforts).

Although our example involved a specific military problem (assessing military capability of alternative military forces to halt an invading army by using long-range fires in the form of aircraft and missiles), we convinced ourselves that the example would illustrate many generic issues.

The base model (called EXHALT-CF)[9] has input variables such as the number of resources always available (forward-deployed shooters, such as fighter aircraft), the rate at which those can be increased (deployment rates), the times at which partial and full rates of increase would be initiated (related to strategic warning, time of decision, time at which access to bases is granted, etc.), and so on—to include the effectiveness of the resources (kills per shooter-day) and the size of the task to be accomplished (the number of threat divisions, etc.). An important output is the distance that would be moved by the attacking army before it is halted. The meta model, we would hope, would be able to predict this distance from a much smaller set of inputs. The inputs could be a subset of the original model's inputs or a set of composite variables such as the sum of two high-resolution inputs (or, realistically, something much more complex).

V. Issues and Hypotheses

Before beginning the experimental phase of our study, we developed a set of issues and hypotheses to guide our exploration. These included the following:

- Black-box models (such as statistical meta models) are less useful to decision makers than phenomenologically motivated models with clear physical interpretations. Thus, if they are to compete effectively, they must be accurate and reliable.
- Statistical meta models may be relatively accurate "on the average," but may be seriously misleading for predicting sensitivities and variation.
- Statistical meta models may be seriously misleading on crucial "system issues" (to be discussed below).
- Some statistical methods may yield expressions with meaningful physical interpretations by "discovering" composite variables.
- The potential advantages of models based in theory (i.e., phenomenological models) may not be realized in practice because the resulting analytical forms turn out to be ugly, complex, and opaque.
- A synthesis of approaches may be desirable: one in which theory is used to guide application of statistical tools.

The first of these reflects our ingoing attitude (statisticians might say bias). In candor, our effort has not really been devoted to finding new statistical methods to improve accuracy. Many first-rate researchers work on such matters and a considerable literature already exists. Instead, our real objective is suggested by the last item in the list: the belief that a synthesis of theory-based and statistical methods might prove practical and attractive. As indicated by the middle items, we also were suspicious about how meta models developed with relatively standard methods—could be on issues of interest to us. Particularly interesting to us here was the “system issue.” By this we mean that many important problems are about assessing the capabilities of systems with multiple individually critical components. Such systems depend for their success on all of these critical components separately proving successful. Not all systems are of this type, but many of interest are. Analytically, to say that a system depends on each of subsystems A, B, and C being successful suggests that overall capability depends on something more like a product of capabilities, $C_A C_B C_C$, than a sum. Figure 4 shows in the representation of a fault tree the structure of the halt problem on which we focused for our example. This fault-tree representation highlights the system character we have in mind: success in achieving an *early* halt of an invasion requires success in each of the four components indicated by branches.

We would not expect normal linear regression to generate good meta models when such system effects are present. Even generalized regression methods, which consider various nonlinear composite variables, typically do not include triplet products. This justified our suspicion, but proved nothing because in practice statistical models often do much better than one would expect a priori. Further, dependences among variables, such as represented by product terms $C_A C_B C_C$ can sometimes be reasonably approximated by a sum of terms such as $C_A C_B$, $C_A C_C$, and $C_B C_D$. We were also impressed by the common lore among statisticians that pair wise interactions among variables are typically sufficient for meta modeling—that diminishing returns sets in quickly in considering interactions. This lore was in conflict with our theory-based reasoning, but merited respect as we constructed hypotheses to explore. Finally, several advanced statistical methods (e.g., cluster methods) appeared to merit investigation if time permitted.

VI. Selected Observations

With this background of motivation and approach, let us now describe briefly some of the observations we have made to date, based on our experiments—which should be viewed more as developing a case history and making observations about it, than as something rigorously systematic.

A. *Success of the Statistical Meta Models*

We ran 1000 cases of our base model, generating them randomly from the input space of the model by representing the input variables with random distributions. We then developed a series of increasingly sophisticated statistical models while avoiding insertion of phenomenology. The meta models were based, in increasing order of sophistication, on:

- Conventional linear regression of all the input variables
- Modestly extended linear regression in which the variables used as the basis for linear regression were composites of the original input variables—composites motivated by looking for consistency of dimensionality in many of the variables regressed. In particular, we constructed a number of composite variables with the dimensions of distance.
- More generalized regression using as the basis not just the original input variables $\{X_i\}$, but also the various product terms $\{X_i X_j\}$.

As expected, the linear regression did not do particularly well (although better than one might expect), but with the embellishments, we obtained fair agreement with the predictions of the actual base model. This conclusion, however, applied only so long as we focused on “standard” measures, such as R^2 or, better, root mean square error. Root mean square error varied from about 60-100 km, depending on which statistical model we attempted. Since the goal was to achieve a halt distance less than 100 km, this degree of variation was not really satisfactory—although, again, it was better than one might expect given the complexity we believed existed in the original model.

When viewed in a more fine-grained way, results were worse. For example, some of the coefficients had nonsensical signs and the errors of individual cases made no sense. But why should they have made sense when the “models” used had little physical content?

Most important, the statistical meta models did not do well when used to compare the relative importance of variables. A basic reason for this is that the statistical meta model is created by reducing average error over the entire input domain. However, in many problem areas—such as military problems where one has a thinking adversary, or an economic domain in which choices are not made randomly but to maximize profit—small “corners” of the input space can be sought out. For example, an adversary may minimize warning time *and* invade rapidly *and* use various tactics to degrade the defense’s capabilities—even if temporarily. Predicting outcomes for a corresponding war might mean running the model for a set of inputs that would be regarded as extremely improbable if they were independent and random. One way to think about this is to refer to the inputs as mathematically independent, but strategically correlated.

It is easy to understand how a purely mathematical effort to assess the relative importance of variables can run into trouble. Such an effort might, for example, measure the average effect of a 1% change in a given variable when averaged over all of the rest of the input space. If that variable was extremely important only in one “corner” of the space, that fact would be lost as the result of the broader averaging.

Another way to think about the problem is to look at graphs comparing predictions of the meta model with the base model. Not uncommonly, the meta model will do poorly in one domain and poorly (but with opposite sign in the error) in another domain. It will also do extremely well in some domains and quite poorly in others, even though, on average, it will do fairly well. When one asks about the validity of an approximation or the relative importance of a variable in such a case, the result will be correct on average but potentially quite misleading.

The problem, some might respond, was in considering too large an input space. In a sense, that is true. However, which "corner" of the space is of interest depends on details of context that are difficult to predict in advance. Nonetheless, this is the essence of the problem.

B. An Infusion of Theory

What happens, then, when we add bits of theory before generating the statistical meta models? Suppose, for example, that a problem has three inputs X,Y, and Z. Adding theory might be to assert that that meta model should have the form $C_1XY/Z + C_2X$. The composite variables forming the dimensions for regression, then, would be Q_1 and Q_2, where $Q_1=XY/Z$ and $Q_2=X$. We have elsewhere called these "aggregation fragments." suggested by theory. Linear regression could then be used to determine the coefficients C_1 and C_2. And, if one were lucky, perhaps C_2 would be small and the meta model could be simply C_1XY/Z.

In more realistic cases, of course, the base model might have dozens of inputs and the composite variables might be complex as well. Further, it might or might not be possible to use linear regression straightforwardly. In the case we worked in detail, for example, the form suggested by theory involved Max and Min operators, which can cause trouble. Tricks can often be applied, however, such as breaking the data into groups and applying the methods of linear regression on the groups separately, or ignoring the Max and Min operators until after finding a regression model and then applying the operators. What is valid depends on details of the problem.

What we learned from our experimental application of our ideas was the following:

- Infusing the approach with theory-motivated aggregation fragments may or may not improve the meta model significantly if the only measure of goodness is something like R^2 or root mean square error.
- However, the resulting meta model will at least have pieces with understandable significance. That is, its descriptive value will be higher.
- Further, the enhanced meta model may be more accurate in predicting relative importances and may help users avoid serious pitfalls. If, for example, one knows that it is the product XY/Z that matters most (although X,Y, and Z may also appear in the definition of some of the less important composite variables), then that could be quite useful in drawing valid conclusions–and ignoring artifactual conclusions—about relative sensitivities. Also, if theory were to tell us that an aggregation fragment $Q_1 = \sum_i^n X_i$ should be important, then one could avoid the error of concluding from a more naive meta model that the individual variables $\{X_i\}$ are unimportant. That is, the coefficients of a naive regression might be only a third as large for each of the X_i, as that for, say, X_{n+1}, but if n were 10, then Q_1 would be more important than X_{n+1}—if only one knew to look for Q_1.
- Most important, perhaps, our experiments confirmed the potential value of imposing a theory-motivated "system structure" on the meta model.

To illustrate this trivially, suppose that we were interested in the rate at which something could be detected from searching an area. Elementary theory would tell us that the rate would depend on the product of search rate R and the probability of detection when viewing an area that in fact contains the item of interest. At a more microscopic level, there might be a great many variables such as the search vehicle's speed, time on station, turnaround time for refueling and repair, search pattern, and so on. Also, the detection probability in the sense that we mean it might not appear. Instead, one might have inputs for the power and aperture of a radar, its scan rate, the radar cross section of interesting objects, the probability of recognizing that a particular moving object was an example of the item in question, and so on. A linear regression of these variables might produce something useful, but would not pick up the right form. If instead the meta model were assumed to have the form RP_d, where R was constructed from the search vehicle's attributes using even something as simple as dimensional analysis, and where P_d was assumed to be a product of the sensor attributes and target cross section (but limited to 1), then the resulting meta model would be guaranteed to have the characteristic that the search would be predicted to be a failure if either R or P_d were too small. That is, one would not make the mistake of predicting that one could compensate for a very poor search platform by upping the performance of the power and aperture of its radar.

In the actual problem that we worked through experimentally, the meta model that we concluded should be tried based on theory had the form shown in the equations below, where the independent variables were Obj (the objective sought by the attacker, corresponding to the distance from his border to a strategically important destination), V (the initial movement rate of the attacker), ξ (the number of attacker armored vehicles that the defender must kill to halt the invasion), δ_{max} (the number of kills each defender shooter can kill each day using the best weapons available), δ_B (the same quantity, but for a poorer weapon available in large numbers), T_{SEAD} (the time required to suppress the attacker's air defenses so that shooters can operate effectively), T_x (the time at which shooters begin their attack on the armored column), R (the rate at which shooters deploy to the region), A_0 (the number of shooters present when the war starts), A_{max} (the maximum number of shooters that can be in the theater), N_{awpn} (the number of top-quality weapons), and Ω (the slowing of the invader's movement for each vehicle killed per day).

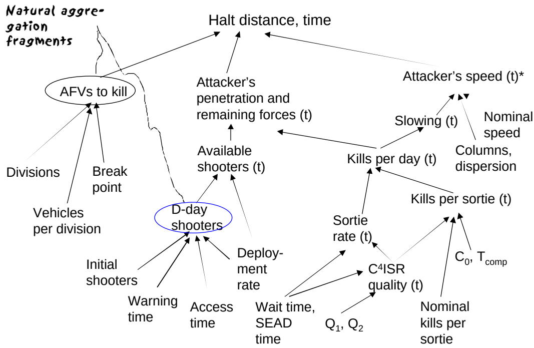

Figure 4—Finding "Aggregation Fragments"

Details are not of interest here, but note that the theory-motivated meta model is quite nonlinear and that it has recognizable "system features" in that, for example, the distance gained by the attacker can be large if it the attacker's size ξ is large or if the defender's per-shooter-day effectiveness δ_{max} or δ_B is low or if the defender has too few shooters on average. The form is not that of a simple product because there are other complications, but that "product" feature is prominent in the expression for the composite variable D_2.

$$D = Max[Min[D_2 - C_1 T_{delay} Obj], 0]$$

$$D_2 = C_2 \frac{\xi_A}{\bar{A}\delta_A} + C_3 \frac{\xi_B}{\bar{A}\delta_B} - C_4 \Omega \xi$$

$$\xi_A = Min[N_{awpns} \frac{\delta}{SN_a}, \xi] \quad \xi_B = \xi - \xi_A$$

$$\bar{A} \approx Min[A_0 + RT_x + \frac{1}{2} R(T_{SEAD} - T_x), A_{\max}]$$

Without elaborating, let it suffice to say that this theory-motivated meta model did spectacularly well—even embarassingly so. We say "embarassing" because the base model took months of work to develop, code, and debug, and is in no way simple and transparent. Nonetheless, the underlying factors driving its results are largely those summarized in the compact expressions above. To someone interested in this particular problem, the structure of this expression and the various terms can be explained clearly in a matter of minutes.

As one would expect, the theory-motivated meta model did well when asked to predict sensitivities and relative importances.

In our experience with this and vaguely similar problems, it has proven possible to develop "smart" suggested meta model forms with hours, days, or a few weeks of work, rather than months. To be sure, this requires shifting mindsets from that often associated with procedural programming to that like more traditional analytical modeling—even with use of paper, pencil, and a whiteboard.

In summary, our experiments tended to confirm the initial hypotheses and to give them sharper meaning. We can hardly draw universal conclusions from such experiments, but we are encouraged that the traditional methods of mathematical modeling and statistical meta modeling can be merged in developing useful low-resolution models that are reasonably suitable for the kind of high-level exploratory analysis needed for both policy planners and certain kinds of intelligent machines.

C. *Other Observations*

Finally, let us comment briefly on some issues that we had found puzzling at the outset. One of these was the common belief among statisticians who generate meta models using experimental designs to sample the results generated by a physical system or base model that interaction effects can typically be ignored beyond those of pairwise interactions. The reason for this is probably just that the applications are limited to problems in which a single nicely behaved "response surface" applies. If that is the case, then—by analogy with Taylor's theorem in calculus—one would expect the quadratic approximation would often be reasonably good. However, in policy problems—including the one that we used for our example—the non linearities caused by thresholds of various kinds result in a more complex and non monotonic structure. No single response surface suffices. Furthermore, in problems with which we are familiar the empirical data or realm of validity for the base model is often quite limited. It is important to be able to extrapolate the meta model's predictions well beyond the region for which it was calibrated. When this is so, it should hardly be surprising that a theory-motivated meta model (perhaps with various If-Then-Else constructions distinguishing broad regions) can be far better than a more naively generated statistical meta model.

VII. Conclusions

In summary, there is great potential in marrying the techniques of statistical meta modeling with the insights of theoretical, phenomenological, modeling. The benefits of such a synthesis are likely to be quite high when attempting to represent systems with individually critical components and complex systems with substantially different behaviors in different regimes of their input variables, and in predicting system behaviors for circumstances significantly different from those for which one has empirical data.

The synthesis we are suggesting rejects the "purist" approach of some statisticians, which is sometimes characterized as "Let the data speak," by which is meant that one should explicitly *avoid* postulating a theoretical structure to the model and instead see what the statistical analysis reveals. Such an approach has much to offer in many problems, but not the ones we are addressing. In our problems, it usually pays to have theory. The payoff is quite high in terms of its cognitive benefits (related to the model's expanatory power), which may be even more important than modest differences in the accuracy or precision of prediction. We believe that will continue to be the case for strategic planning. It may or may not be true in the long run for robots in cases where the data available for calibrating a meta model is massive and credible, but we suspect that paucity and unreliability of data will plague intelligent systems used in complex environments (e.g., planetary explorers rather than spot welders).

In attempting a synthesis of approaches, we suggest several principles:

- Attempt to characterize the problem using the methods of multiresolution, mutiperspective modeling (MRMPM)—especially the method of hierarchical or nearly hierarchical decomposition.
- Attempt to find meaningful simplified structures by sharpening the hierarchicies—i.e., by identifying approximations (perhaps case-dependent approximations) that create nearly decomposable hierarchies.
- In doing so, however, be guided less by the intuition or preferences of pure mathematics (e.g., independent events) than by the character of the actual problem. Worry about what we have called "strategic correlations."
- Attempt to characterize the problem "formally" even if one cannot as a practical matter accomplish the various computations implied. Attempt to structure the problem so as to "see" system features where one knows they should exist, but allow structurally for complications (e.g., even if unusual, it may be possible for one component—if present in quantity–to substitute for another thought to be individually critical).
- Abstract from this theoretical work both "aggregation fragments" and structure that can be used to inform statistical meta modeling.
- Try to identify variables that are being short-changed in the proposed structure and then avoid using the meta model for predicting the consequences of change in those variables, even though the meta model depends on them.

We are nowhere near providing firm principles or recipes for success, but we believe that the approach we suggest will prove quite useful. One reason for our belief here is that the suggestions appear to be in some respects a restatement—for a new context of inquiry—of methods that have long been applied by physical scientists and engineers.

ACKNOWLEDGMENT

This paper is based on research accomplished for the Air Force Research Laboratory. It also draws on work accomplished for the Office of the Secretary of Defense.

REFERENCES

[1] National Institute of Standards and Technology, Measuring Performance of Systems with Autonomy: a White Paper for Explaining Goals of the Workshop, final draft, Alex Mystel, July, 2000.

[2] Meystel, Alex, "Evolution of Intelligent Systems Architectures: What Should Be Measured," PERMIS 2000 proceedings.

[3] Meystel, Alex and Elena Messina (ed), *Measuring the Performance and Intelligence of Systems: Proceedings of PerMIS 2000*, National Institute of Standards and Technology, Gaithersburg, 2001.

[4] Messina, Elena, Alex Meystel, and Larry Reeker, "PERMIS' 2001 White Paper: Measuring Performance and Intelligence of Intelligent systems,"PERMIS' 2001, Mexico City, 2001.

[5] Simon, Herbert, *Sciences of the Artificial*, MIT Press, 3d. ed., Cambridge, MA, 1996.

[6] Davis, Paul K. and Richard Hillestad, "Families of Models That Cross Levels of Resolution: Issues for Design, Calibration and Management," *Proceedings of the Winter Simulation Conference*, 1992.

[7] Davis, Paul K. and James H. Bigelow, Experiments in Multiresolution Modeling, RAND MR-1004-DARPA, Santa Monica, CA, 1998.

[8] Davis, Paul K., James H. Bigelow, and Jimmie McEver, Effects of Terrain, Maneuver Tactics, and C4ISR on the Effectiveness of Long-Range Precision Fires: a Stochastic Multiresolution Model (PEM) Calibrated to High-Resolution Simulation MR-1138-OSD, RAND, Santa Monica, CA, 2000[2]Meystel, Alex, "Evolution of Intelligent Systems Architectures: What Should Be Measured," *PerMIS 2000 Proceedings*, National institute of Standards and Technology, August, 2000.

[9] Davis, Paul K., Jimmie McEver and Barry Wilson, *Measuring Interdiction Capabilities in the Presence of Anti-Access Strategies*, RAND MR-1471-AF, forthcoming.