

Intelligent Model Selection for Hillclimbing Search in Computer-Aided Design

Thomas Ellman John Keane
Department of Computer Science
Rutgers University
{ellman,keane}@cs.rutgers.edu

Abstract

Several investigators have recently developed methods of automatically selecting among multiple models of physical systems. Our research is novel in that we are developing model selection techniques specifically suited to computer-aided design. Our approach is based on the idea that artifact performance models for computer-aided design should be chosen *in light of the design decisions they are required to support*. We have developed a technique called "Gradient Magnitude Model Selection" (GMMS), which embodies this principle. GMMS operates in the context of a hillclimbing search process. It selects the simplest model that meets the needs of the hillclimbing algorithm in which it operates. We are using the domain of sailing yacht design as a testbed for this research. We have implemented GMMS and used it in hillclimbing search to decide between a computationally expensive potential-flow program and an algebraic approximation to analyze the performance of sailing yachts. Experimental tests show GMMS to achieve a tradeoff between the quality of a yacht design and the amount of time needed to find it.

1. Introduction

Models of a given physical system can differ along several dimensions, including the cost of using the model, the accuracy and precision of the results, the scope of applicability of the model and the data required to execute the model, among others. More than one model is often needed because different tasks require different tradeoffs among these dimensions. A variety of criteria and techniques have been proposed for selecting among various alternative models of physical systems. Several investigators have recently developed methods of automatically selecting among multiple models of physical systems. For example, some techniques select appropriate models by analyzing the struc-

ture of the query the model is intended to answer [Falkenhainer and Forbus, 1991], [Ling and Steinberg, 1992], [Weld and Addanki, 1991]. Another approach selects an appropriate model by reasoning about the simplifying assumptions underlying the available models [Addanki *et al.*, 1991]. Yet another approach reasons about the accuracy of the results the model is required to produce [Weld, 1991]. The importance of this subject is evidenced by the number of papers on model selection appearing in two recent workshops on the subject [Ellman, 1990], [Ellman, 1992]. Our research is novel in that we are developing model selection techniques specifically suited to computer-aided design. Our approach is based on the idea that artifact performance models for computer-aided design should be chosen *in light of the design decisions they are required to support*. We have developed a technique called "Gradient Magnitude Model Selection" (GMMS), which embodies this principle. GMMS operates in the context of a hillclimbing search process. It selects the simplest model that meets the needs of the hillclimbing algorithm in which it operates.

The GMMS technique is being developed and tested in the context of the "Design Associate", an interactive environment for designing complex engineering structures [Ellman *et al.*, 1992]. The domain of sailing yacht design is being used as a testbed for this research. An example of a sailing yacht, the Stars and Stripes '87, is shown in Figure 1. This yacht figures prominently in the history of yacht racing. In particular, it won the America's Cup back from Australia in 1987 [Letcher *et al.*, 1987]. The Design Associate is being developed with an eye toward reconstructing the design process that led to the Stars and Stripes '87 yacht.

Racing yachts can be designed to meet a variety of objectives. Possible yacht design goals include: *Course Time Constraints*, *Rating Constraints* and *Cost Constraints*. These goals are summarized in Figure 2. A course time goal is specified by three pieces of information: A race course description, an expected wind speed and an upper bound on the

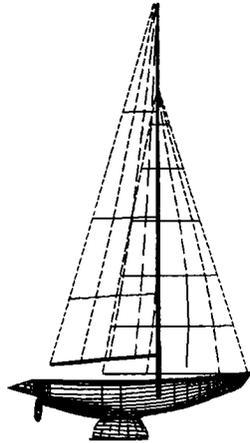


Figure 1: The Stars and Stripes '87

$$\begin{aligned} \text{CourseTime}(\text{Yacht}, \text{Course}, \text{WindSpeed}) &\leq T \\ \text{Rating}(\text{Yacht}) &\leq R \\ \text{Cost}(\text{Yacht}) &\leq C \end{aligned}$$

Figure 2: Racing Yacht Design Goals

time for the yacht to traverse the given course under the given wind conditions. In the Design Associate, the course time is estimated using a velocity prediction program, called "VPP" [Letcher, 1991]. The rating goal is imposed by the authorities who manage the America's Cup races. It requires that the "rating" of each yacht be at most "12 meters". Ratings of yachts are determined by a complex set of rules known as the "International Measurement System" devised by the Offshore Racing Council [Council, 1991]. The Design Associate estimates the rating of a yacht using a simplified formula depending on a few major geometric measurements of the yacht hull. Finally, a cost goal is specified by providing an upper bound on the cost of building the yacht. The Design Associate estimates cost using another simple formula depending on a few major geometric measurements of the yacht hull.

Of all the quantities appearing in the design goals, course time is the most difficult to evaluate. The VPP program that performs this task is described in Figure 3. VPP takes as input a set of B-Spline surfaces representing the geometry of the yacht hull [Rogers and Adams, 1990]. Each surface is itself represented as a matrix of "control points" that define its shape. VPP begins by using the "hull processing models" to determine physically meaningful quantities impacting on the performance of the yacht, e.g., wave resistance (R_w), friction resistance (R_f), effective draft (T_{eff}), vertical center of gravity (V_{cg}) and vertical center of

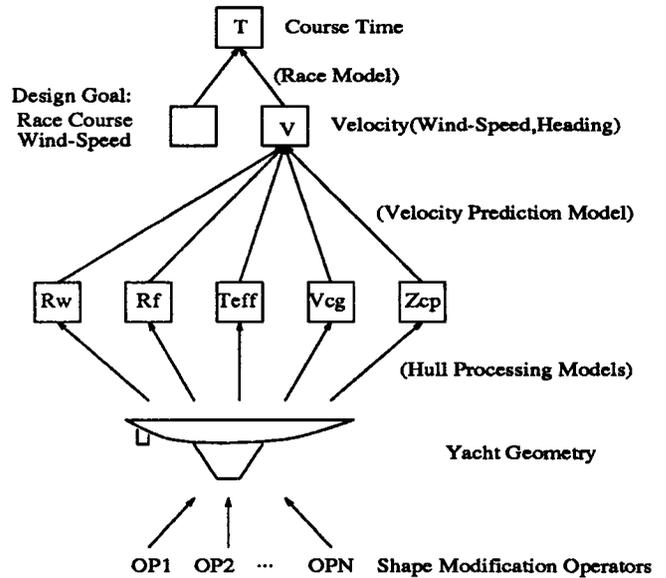


Figure 3: Velocity Prediction Program

pressure (Z_{cp}), among others. These quantities are then used in the "velocity prediction model" to set up non-linear equations describing the balance of forces and torques on the yacht. The velocity prediction model uses an iterative method to solve these equations and thereby determine the "velocity polar", i.e., a table giving the velocity of the yacht under various wind speeds and directions of heading. Finally, the "race model" uses the velocity polar to determine the total time to traverse the given course, assuming the given wind speed.

2. Modeling Choices in Yacht Design

A number of modeling choices arise in the context of sailing yacht design. These choices are outlined in Figure 4. Probably the most important is the choice of models for estimating the effective draft (T_{eff}) of a yacht. Effective draft is a measure of the amount of drag produced by the keel as a result of the lift it generates. An accurate estimate of this quantity is quite important for analyzing the performance of a sailing yacht. Unfortunately, the most accurate way to estimate effective draft is to run a highly expensive potential flow code called PMARC. (This code takes approximately one hour when running on a Sun Microsystems Sparc 2 Workstation.) Effective draft can also be estimated using an algebraic approximation with the general form outlined below:

$$\begin{aligned} T_{eff} &= K \sqrt{D^2 - 2A_{ms}/\pi} \\ D &= \text{Maximum Keel Draft} \\ A_{ms} &= \text{Midship Hull Cross Section Area} \end{aligned}$$

- Algebraic Approximations v. Computational Fluid-Dynamics: The effective draft T_{eff} of a yacht can be estimated using an algebraic approximation or by using a potential flow code called "PMARC".
- Reuse of Prior Results v. Recomputation of Results: Some physical quantities may not change significantly when a design is modified. For a given physical quantity, its value may be retrieved from a prior candidate design, or its value may be recomputed from scratch.
- Linear Approximations v. Non-Linear Models: Velocity polars can be computed as linear functions of resistances and geometric quantities or by directly solving non-linear force and torque balancing equations.

Figure 4: Modeling Choices in Yacht Design

This algebraic model is based on an approximation that treats a sailing yacht hull as an infinitely long cylinder and treats the keel of the yacht as an infinitely thin fin protruding from the cylinder. The constant K is chosen to fit the algebraic model to data obtained from wave tank tests, or from sample runs using the PMARC potential flow code. Although the algebraic approximation is comparatively easy to use, its results are not as accurate as those produced by the PMARC potential flow code.

Another important modeling choice involves the decision of when to reuse the results of a prior computation. The importance of this type of decision is illustrated by Figure 5. Suppose a designer is systematically exploring combinations of canoe-bodies and keels of a sailing yacht. In order to evaluate the performance of a yacht, he must evaluate the yacht's wave resistance R_w as well as its effective draft T_{eff} . Wave resistance depends mainly on the canoe-body of the yacht and is not significantly influenced by the keel. Thus when only the keel is modified, wave resistance will not significantly change. Thus, instead of recomputing wave resistance for the new yacht, the system can reuse the prior value. On the other hand, effective draft depends mainly on the keel of the yacht and is not significantly influenced by the canoe-body. Thus when only the canoe-body is modified, effective draft will not significantly change. Thus, instead of recomputing effective draft for the new yacht, the system can reuse the prior value. In fact, the entire matrix of yachts can be evaluated by computing wave resistance for a single row, and computing effective draft for a single column. Thus by intelligently deciding when to reuse prior evaluation results, the system can achieve considerable savings in the computational costs of design.

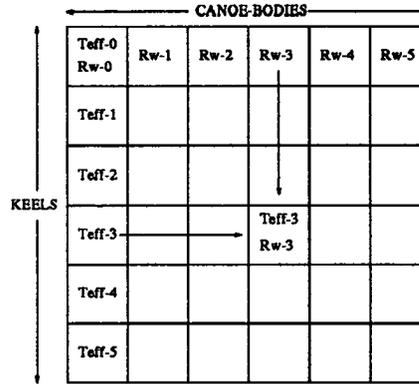


Figure 5: Reuse of Prior Results

3. Gradient Magnitude Model Selection

Gradient Magnitude Model Selection (GMMS) is a technique used in the Design Associate for selecting evaluation models in the context of a hillclimbing search procedure. The key idea behind this technique is illustrated by Figure 6. Suppose the system is running a hillclimbing algorithm to minimize *CourseTime* as estimated by some approximate model. The values of *CourseTime* returned by this approximate model are indicated by the curved line. Suppose further that the system is considering the hillclimbing step illustrated in the figure. If the error bars shown with solid lines reflect the uncertainty of the approximate model, the system can be sure that the proposed step will diminish the value of *CourseTime*. On the other hand, using the error bars shown with dotted lines, the system would be uncertain as to whether the true value of *CourseTime* would improve after taking the proposed hillclimbing step. In the first case, the system could safely use the approximate model to decide whether to take the proposed hillclimbing step, while in the second case, the approximate model would not be safe to use for that decision. Thus GMMS evaluates the suitability of an approximate model by comparing error estimates to the magnitude of the change in the optimization criterion as measured by the approximate model. In order to operate, GMMS requires an approximate model, a "true" model, and some means of estimating the error of the approximate model in comparison to the "true" model. GMMS can in principle be applied to any of the modeling choices outlined in Figure 4.

In our research, we have experimented with GMMS using the choice of models for effective draft, T_{eff} , as a test case. Thus GMMS chooses between the algebraic approximate model and the PMARC potential flow model described above. Our system obtains an error estimate for the al-

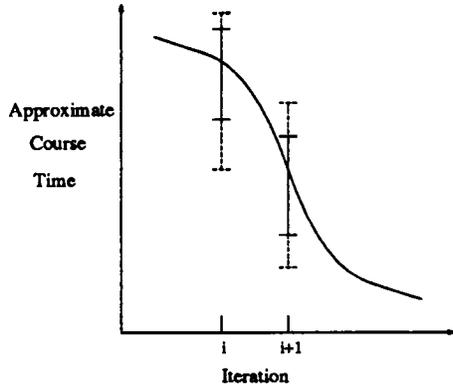


Figure 6: Gradient Magnitude Model Selection

gebraic model using the procedure outlined in Figure 7. The procedure takes as input two sets, A and B , of sample points in the space of candidate yacht designs. The set A is a small, sparsely distributed point set, while set B is a larger, more densely distributed point set. The system constructs two versions of the algebraic model, by choosing values for the fitting coefficient K , described above. $Alg(A)$ is fitted against the “true” values from the sparse point set A . $Alg(B)$ is fitted against the “true” values from the dense point set B . In each case the “true” values are determined using the *PMARC* as the “gold standard”. Since $Alg(B)$ is fitted against the denser point set, this model is actually used during hillclimbing search; however, its error is actually estimated using $Alg(A)$, which was fitted against the sparser point set. In particular, the error in $Alg(B)$ is estimated by comparing $Alg(A)$ to *PMARC* for all points in the set $B - A$. Two different error estimates, result from this procedure: *Absolute-Error* is based on the assumption that errors in the algebraic model at nearby points in the design space are independent of each other. *Difference-Error* takes into account the possibility that errors for nearby points may be correlated.

Gradient-magnitude model selection (GMMS) operates in a slightly different manner depending on which type of error estimate is available. Consider first how *Absolute-Error* estimates are used. Given two candidate yacht designs D_i and D_{i+1} , the system first evaluates the effective draft T_{eff} of each candidate using the algebraic approximation. The estimate of *Absolute-Error* is then used to find upper and lower bounds on the effective draft T_{eff} of each candidate. Each pair of bounds is then propagated through the rest of the velocity prediction program (Figure 3) to obtain an upper and lower bound on the *CourseTime* of each candidate. If the *CourseTime* intervals do not overlap, then the system knows whether to take a step from candidate D_i to D_{i+1} based on the results of the algebraic approximate model. On the other hand, if

1. Let A be a sparse point set in the design space (u_1, \dots, u_n) .
 - (a) Run *PMARC* to find $T_{eff}(u)$ for each point u in the set A .
 - (b) Fit coefficients in $Alg(A)$ to minimize average error over the set A .
2. Let B be a sparse point set in the design space (u_1, \dots, u_n) .
 - (a) Run *PMARC* to find $T_{eff}(u)$ for each point u in the set B .
 - (b) Fit coefficients of $Alg(B)$ to minimize average error over the set B .
3. Estimate the error of $Alg(A)$ using the *PMARC* as the “gold standard”:
 - *Absolute-Error*($Alg(A)$) = Average error in T_{eff} over all points in $B - A$.
 - *Difference-Error*($Alg(A)$) = Average error in $|T_{eff}(u) - T_{eff}(u')|$ as function of $\Delta u_1, \dots, \Delta u_n$, over all pairs (u, u') of points in $B - A$.

Figure 7: Error Estimation Technique

the intervals do overlap, then the system must use *PMARC* to obtain a better estimate of effective draft T_{eff} for each candidate. When *Difference-Error* estimates are available, GMMS operates differently. After computing the effective draft of each candidate, the system considers two scenarios: (1) All of the *Difference-Error* occurs in the T_{eff} of D_i , and none occurs in the T_{eff} of D_{i+1} ; (2) All of the *Difference-Error* occurs in the T_{eff} of D_{i+1} , and none occurs in the T_{eff} of D_i . In each case the system propagates the *Difference-Error* through the rest of the velocity prediction program, to obtain bounds on the *CourseTime* of each candidate. The algebraic approximate model is considered acceptable only if the *CourseTime* intervals are disjoint under both scenarios. Similar methods can be used to apply Gradient Magnitude Model Selection to the other modeling choices shown in Figure 4.

4. Experimental Tests

Results from a preliminary test of gradient magnitude model selection are show in Figure 9. The data in this graph was collected by using a hillclimbing algorithm to modify the Stars and Stripes '87 yacht to sail in a new environment with an average wind speed of 16 knots. The graph shows how the value of *CourseTime* changed during the hillclimbing process. Points marked with a diamond indicate *CourseTime* values computed using the algebraic approximation. Points marked with a plus sign indicate *CourseTime* values computed using the *PMARC* potential flow code. (The

	Course Time	PMARCs
PMARC Only	11730.8	281
GMMS	11747.0	104
Algebraic Only	11757.9	0

Figure 8: Performance of Gradient Magnitude Model Selection

graph shows only a fraction of the number of evaluations carried out during the entire design process. In particular, evaluations used to compute gradients are not shown.) Notice that the algebraic model was sufficient to verify progress toward a better design during iterations zero through five. During this portion of the search, the changes in *CourseTime* at each iteration were quite large. Improvements could therefore be verified using the algebraic model. During subsequent iterations, the *CourseTime* curve leveled off. The system was then forced to switch to the *PMARC* model in order to verify improvements from one iteration to the next.

Additional experiments have been run in order to demonstrate the practical value of gradient magnitude model selection (GMMS). Results from a comparison of three hillclimbing search runs are shown in Figure 8. These three runs include: (1) A run using *PMARC* to compute effective draft; (2) A run using GMMS to choose between *PMARC* and the algebraic approximation; and (3) A run using only the algebraic approximation. Each run attempted to redesign the Stars and Stripes '87 yacht to sail in a new environment with an average wind speed of 16 knots. The three runs are compared according to the *quality* of the solution, measured in terms of the *CourseTime* of the yacht, and the *design time* needed to obtain the solution, measured in terms of the number of *PMARC* evaluations needed to obtain the solution. (*PMARC* is, by far, the most expensive of the velocity prediction program shown in Figure 3.) Notice that the GMMS run found a yacht that was better in quality (lower in *CourseTime*) than the one found using the algebraic approximation, yet GMMS required less design time (fewer *PMARC* evaluations) than would have been necessary had *PMARC* been used all the time. Thus GMMS implements a *tradeoff* between the quality of design and the design time needed to find it. In the future we plan additional tests of the practical value of GMMS. In particular, we plan to test the value of GMMS when used to make additional modeling choices, such as those listed in Figure 4, that arise in the context of racing yacht design.

5. Acknowledgments

The research reported in this paper has benefited from numerous discussions with all members of the

Rutgers CAP project, including Saul Amarel, Andrew Gelsey, Haym Hirsh, Chun Liew, Ringo Ling, Gerry Richter, Louis Steinberg, Mark Schwabacher, Chris Tong and Ke-Thia Yao. This research was supported by the Defense Advanced Research Projects Agency (DARPA) and the National Aeronautics and Space Administration (NASA), under DARPA-funded NASA grant No. NAG 2-645.

References

- S. Addanki, R. Cremonini, and J. Scot. Graphs of models. *Artificial Intelligence*, 50, 1991.
- Offshore Racing Council. *International Measurement System: A Handicapping System for Cruising/Racing Yachts*. Offshore Racing Council, Newport, Rhode Island, 1991.
- T. Ellman, J. Keane, and M. Schwabacher. The Rutgers cap project design associate. Technical Report CAP-TR-6, Department of Computer Science, Rutgers University, New Brunswick, NJ, 1992.
- T. P. Ellman, editor. *Working Notes of the AAI Workshop on Automatic Generation of Approximations and Abstractions*. Department of Computer Science, Rutgers University, New Brunswick, New Jersey, 1990.
- T. P. Ellman, editor. *Working Notes of the AAI Workshop on Approximation and Abstraction of Computational Theories*. Department of Computer Science, Rutgers University, New Brunswick, New Jersey, 1992.
- B. Falkenhainer and K. Forbus. Compositional modeling: Finding the right model for the job. *Artificial Intelligence*, 50, 1991.
- J. S. Letcher, J. K. Marshall, J. C. Oliver, and N. Salvesen. Stars and stripes. *Scientific American*, 257(2), August 1987.
- J. S. Letcher. *The Aero/Hydro VPP Manual*. Aerh/Hydro, Incorporated, Southwest Harbor, Maine, 1991.
- R. Ling and L. Steinberg. Model generation from first principles. March 1992.
- David F. Rogers and J. Alan Adams. *Mathematical elements for computer graphics*. McGraw-Hill, 2nd edition, 1990.
- D. Weld and S. Addanki. Query-directed approximation. Technical Report 90-12-02, Department of Computer Science and Engineering, University of Washington, Seattle, WA, 1991.
- D. Weld. Reasoning about model accuracy. Technical Report 91-05-02, Department of Computer Science and Engineering, University of Washington, Seattle, WA, 1991.

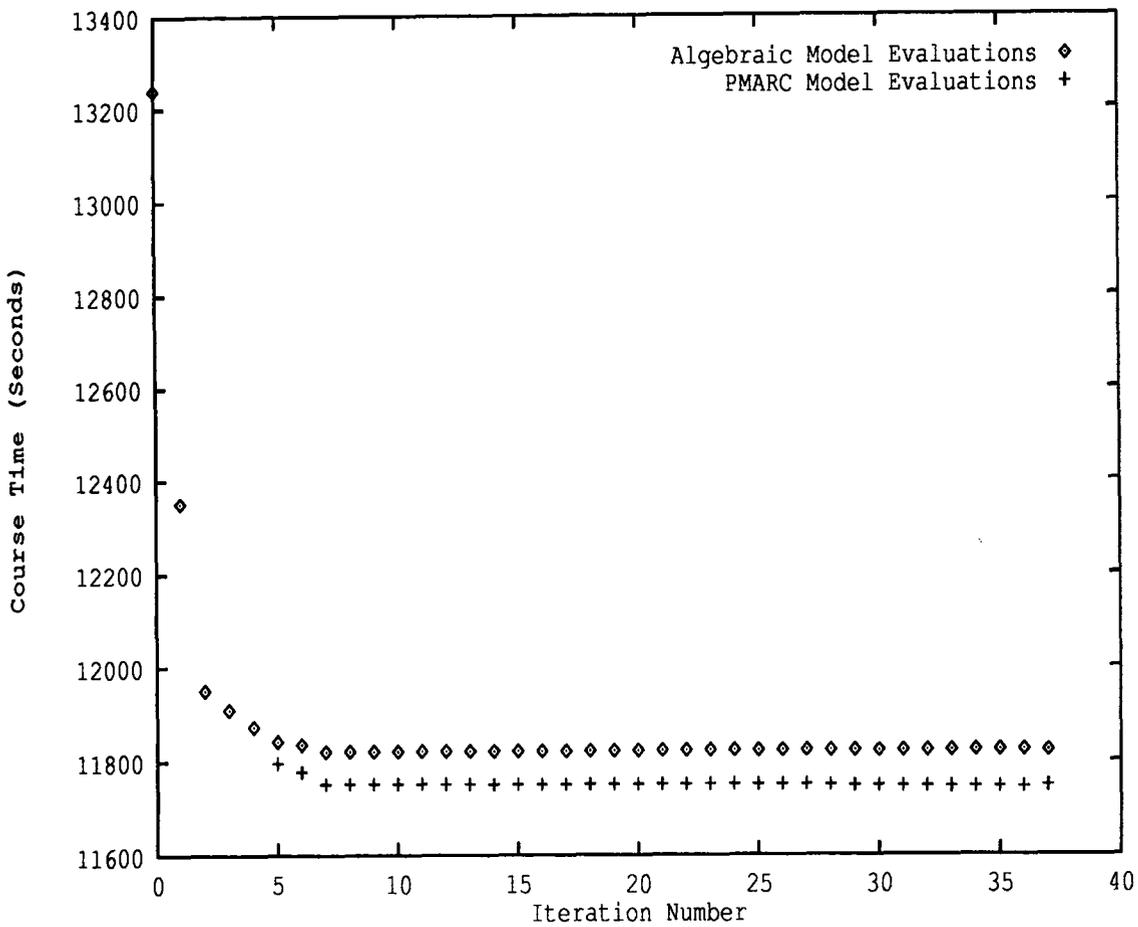


Figure 9: Example of Gradient Magnitude Model Selection

Research Summary

Thomas Ellman John Keane Mark Schwabacher
Department of Computer Science, Rutgers University
New Brunswick, New Jersey 08903 (908) 932-4184
{ellman,keane,schwabac}@cs.rutgers.edu

1. The CAP Project Design Associate

The Design Associate is intended to be an interactive environment that supports decision-making, performance evaluation, design-record management and knowledge acquisition tasks in computer-aided design. The system is specifically intended to handle the design of complex, physical structures, such as ships and planes, among others. The domain of sailing yacht design is being used as a testbed for developing the design associate. Two key difficulties characterize this type of design problem: (1) Design goals depend on global properties of an artifact. (2) Evaluation of the performance of an artifact is computationally expensive. The Design Associate provides a set of tools for attacking each of these problems: Global constraints are attacked by methods of automatically abstracting and decomposing search spaces. Computational costs of evaluation are attacked by methods of intelligently selecting evaluation models at varying levels of approximation. Future work will extend the Design Associate by building a compiler for automatically generating new performance evaluation models. This extension is intended to support innovative design by diminishing the time and monetary costs of developing new models needed to evaluate radically new designs. Research on the Design Associate is expected to contribute to the field of Computer-Aided Design by formalizing the generic task structure of complex, physical structure design. It is also expected to contribute to the field of Artificial Intelligence, by attacking problems such as search control, search space formulation, abstraction, decomposition, model selection and model formation. (See [Ellman *et al.*, 1992].)

2. Abstraction and Decomposition

One portion of our research is focused on automatic decomposition of design optimization problems. Decomposition is especially important in the design of complex physical shapes such as yacht hulls. Exhaustive optimization is impossible because hull shapes are specified by a large number of parameters. Decomposition diminishes optimization costs by partitioning the shape parameters into non-interacting or weakly-interacting sets. We have developed a combination of empirical and knowledge-based techniques for finding useful decompositions. The knowledge-based method examines a declarative description of the function to be optimized in order to identify parameters that potentially interact with each other. The empirical method runs computational experiments in order to determine which potential interactions actually do occur in practice. We

expect this approach to find decompositions that will result in faster optimization, with a minimal sacrifice in the quality of the resulting design. Implementation and testing of this approach are currently in progress.

3. Model Selection

Another portion of our research is focused on intelligent model selection in design optimization. The model selection problem results from the difficulty of using exact models to analyze the performance of candidate designs. For example, in the domain of racing yacht design, an exact analysis of a yacht's performance would require a computationally expensive solution of the Navier-Stokes equations. Approximate models are therefore needed in order to diminish the costs of analyzing and evaluating candidate designs. In many situations, more than one approximate model is available. For example, in the yacht design domain, the induced resistance of a yacht can be predicted by solving La Place's equation - an approximation of Navier-Stokes - or by using a simple algebraic formula. The two approximations differ widely in both the costs of computation and the accuracy of the results. Intelligent model selection techniques are therefore needed to determine which approximation is appropriate during a given phase of the design process.

We have attacked the model selection problem in the context of hillclimbing optimization. We have developed a technique which we call "gradient magnitude model selection". This technique is based on the observation that a highly approximate model will often suffice when climbing a steep slope, because the correct direction of change is easy to determine. On the other hand, a more accurate model will often be required when climbing a gradual incline, because the correct direction of change is harder to determine. Our technique operates by comparing the estimated error of an approximation to the magnitude of the local gradient of the function to be optimized. An approximation is considered acceptable as long as the gradient is large enough, or the error is small enough, so that each proposed hillclimbing step is guaranteed to improve the value of the goal function. Implementation and testing of this approach are currently in progress.

References

T. Ellman, J. Keane, and M. Schwabacher. The Rutgers cap project design associate. Technical Report CAP-TR-6, Department of Computer Science, Rutgers University, New Brunswick, NJ, 1992.