



International Journal on Artificial Intelligence Tools  
© World Scientific Publishing Company

## Learning Satisficing Control Policies for Software Projects

Tim Menzies

*Lane Department of Computer Science and Electrical Engineering, West Virginia University,  
Morgantown, WV, 26501  
tim@menzies.us*

James Kiper

*Computer Science and Systems Analysis Department, School of Engineering & Applied Science,  
Miami University  
kiperjd@muohio.edu*

Jeremy Greenwald

*Computer Science, Portland State University  
jegreen@cecs.pdx.edu*

Ying Hu

*Software designer in Vancouver, British Columbia  
huying\_@yahoo.com*

David Raffo

*School of Business Administration, Portland State University  
raffod@sba.pdx.edu*

Siri-on Setamanit

*Portland State University  
sirion@pdx.edu*

Received (Day Month Year)

Revised (Day Month Year)

Accepted (Day Month Year)

Models in software engineering allow developers to define, but not necessarily explore, the space of design options. Building a useful model, understanding all its interactions, can be intellectually difficult- particularly for large non-linear discrete models. We show theoretically and empirically that our TAR2 *minimal contrast set* learner can generate very succinct conclusions from seemingly complex spaces. Treatment learning is indicated for high-dimensionality non-linear, non-continuous models built in domains with noise or other uncertainties; where managers do not have full control over all model inputs; where there is a requirement to explain or audit models; when reducing model variance is valuable, and there does not exist a budget for building and using maximal models.

*Keywords:* model-driven software engineering, contrast set learning, treatment learning

## 1. Introduction

Consider a software manager using a software process model to control their project. Exploring all the “what-ifs” within those models can be a daunting task. For example, 20 binary decisions require, in the worst case,  $2^{20} > 1,000,000$  what-ifs. After interviewing the managers of three NASA systems regarding their process options, we found that they were exploring up to  $10^9$  options<sup>a</sup>.

In the 21<sup>st</sup>, it is possible to use standard desktop machines to explore a very large option space. Many businesses under-utilize their CPUs once staff go home for the day. Those machines can be combined into a CPU farm which. If the process models are not database or GUI intensive, then such a CPU farm can work through billions of options in an overnight run.

However, once all those options are created, they must be summarized. That is, it is not enough to *generate* a multitude of alternatives, we must also *select* the subset of the possible options that matter the most. Traditional methods may not yield such succinct solutions. For example, in *binary plotting*,  $N$  model inputs are plotted against some output score in  $N$  two-dimensional binary plots. The best constraints on the inputs are those that select for the best output. However, if noise in a model yields plots that are not smooth, then it becomes a time-consuming and subjective task to examine  $N$  binary plots to detect regions of input parameters that most improve the outputs. Also, binary plotting will not find *conjunctions* of input constraints that contribute to better output. Further, as the dimensionality of model input increases, the cognitive effort required to study the plots can become prohibitively expensive.

Many researchers have developed impressive visual environments for decreasing the cognitive overload associated with exploring a multi-dimensional space. Such visualizations help analysts explore visual information, but they present their own challenges. Nevertheless, there are still limits on how many dimensions can be displayed (e.g. we have yet to see effective visualizations for more than a ten dimensionality space). As the visualization environment grows more sophisticated, some users find they have traded a data browsing problem with the new problem of exploring the full range of the effects of all the controls.

When manual exploring of options fail, automatic methods can be applied. Optimization packages can be applied to data or the equations of a system to find “sweet spots” that maximize the score resulting from model outputs. For example, a canonical *sensitivity analysis* method [1] might be to compute eigenvectors of a linear system in order to understand its long-term temporal behavior [2, 3]. Alternatively, using design of experiments, we might exercise an existing model to shed light on the response surface of the model (DOE does identify gradients and key parameters for a model).

<sup>a</sup>A COCOMO model with 12 unknowns was being analyzed; each unknown has up to 6 possible values;  $6^{12} > 1,000,000,000$ .

While useful for some models, these automatic methods do not apply to all models. For example, optimization methods can fail for non-linear models. Any model with an “if” statement introduces a “cliff” where the effects of inputs on outputs abruptly changes. For such models, there is no linear continuous solution that applies either side of the “cliff”. Model uncertainty also complicates sensitivity analysis. For example, the eigenvector technique described would yield spurious results if the coefficients on the models are not known with certainty. Lastly, a DOE analysis can be complicated by the dimensionality, noise, and visualization problems described above.

More generally, traditional methods tacitly assume total knowledge and total control of the input space. This is often impractical. Consider, for example, Mintzberg [4] study of how real-world decisions are made. He found:

- 56 U.S. foremen who averaged 583 activities in an eight-hour shift (one every 48 seconds).
- 160 British managers who worked for half an hour or more without interruption only once every two days.

These empirical observations do not fit a traditional optimization model of decision making were all options are systematically studied, organized, co-ordinated and controlled. In *hard modeling problems*, human agents must make decisions using:

- limited time and computational ability;
- limited computational ability (or limited time for computation);
- limited knowledge about decision alternatives;
- uncertainty about possible outcomes of decisions
- uncertainty about pay-offs;
- no more than a partial ordering of preferences;
- limited information about probabilities of outcomes.

Herbert Simon [5] defined and explored such hard modeling problems using a data structure called *state space* [6]. In terms of model-based SE, state space is the set of options and option selection operators within a model. Further, in hard modeling problems, large portions of the state space are uncertain. Simon argued that in such state spaces, searching for optimal solutions is a spurious goal. Rather, agents can only make just enough decisions that are just good enough. In Simon’s terminology, such decisions are *satisficing*.

Our contribution to hard modeling is to comment that (1) *satisficing* solutions can be achieved by ignoring certain irrelevant or redundant details within a model; (2) surprisingly simple methods can find what irrelevancies to ignore. Treatment learners propose “treatments”; i.e. constraints on a small subset of the model inputs. The other inputs are left to vary at random. That is, apart from generating very succinct solutions, treatment learning offers solutions that are stable despite uncertainty in the non-treated variables.

The rest of this paper is structured as follows. The twin topics of model-based SE and search-based SE are introduced. A special case of *hard models* are discussed that are difficult to manage using standard methods. Evidence is then presented that, often, the key variables that control a model are few in number. The problem of hard models therefore can be reduced to just the problem of understanding the key variables. This will motivate the design of the TAR2 *treatment learner*. Case studies will then be presented showing that, in many domains, TAR2 finds a small number of controlling variables. Such succinct controllers have many advantages:

- Smaller models are easier to explain (or audit).
- Miller has shown that models generally containing fewer variables have less variance in their outputs [7].
- The smaller the model, the fewer are the demands on interfaces (sensors and actuators) to the external environment. Hence, systems designed around small models are easier to use (less to do) and cheaper to build.

In short, for a certain class of *hard modeling* problems, the TAR2 treatment learner can dramatically simplify the task of understanding the configuration possibilities within model-based SE. Hence, our conclusion will be that it is useful to augment standard modeling methods with treatment learning. As to other modeling problems, TAR2 is *not* indicated for low-dimensionality linear continuous models built in domains that have no noise or other uncertainties, and where managers have full control over all model inputs.

## 2. Modeling and Search

Anthropologists argue that the ability to build abstract models is what gives homo sapiens their competitive edge. In his article “What Models Mean” [8], Seidewitz describes the interactions and relationships among the concepts of a model. He asserts that “a model’s meaning has two aspects: the model’s relationship to what’s being modeled and to other models derivable from it. Carefully considering both aspects can help us to understand how to use models to reason about the systems we build ...” These are the aspects of a model’s meaning on which we focus in *model-based* software engineering.

Model-based software engineering (SE) is becoming increasingly important. Sendall and Kozacaynski argue that increasing productivity and reduced time-to-market for software products can accrue when “using concepts closer to the problem domain ...” via modeling [9]. Hailpern and Tarr observe that model-driven development “imposes structure and common vocabularies so that artifacts are useful for their main purpose in their particular stage in the life cycle.” [10]

The utility of model-based SE is widely acknowledged. For example, the Object Management Group (OMG) has recently adopted a *model-driven architecture* framework with goals of “portability, interoperability and reusability through architectural separation of concerns.” [11] Also, Microsoft has been developing a *Software*

*Factory* concept that leverages models with the goal of automation of the software development process [12]. And, at Lockheed Martin, engineers have developed an integrated modeling method called *Model Centric Software Development* that uses “automated generation of partial implementation artifacts,” reverse engineering to integrate legacy assets, and model verification and checking [13]

The importance of models and the model-driven approach is not limited to software design and UML. Many tools exist for modeling such as distributed agent-based simulations [14], discrete-event simulation [15–19], continuous simulation (also called system dynamics) [20, 21], state-based simulation (which includes petri net and data flow approaches) [22–24], hybrid-simulation (that combines discrete event simulation and systems dynamics) [25–27], logic-based and qualitative-based methods [28, chapter 20] [29], and rule-based simulations [30]. One can find models used in the requirements phase (see the DDP method and tool [31] which are a risk-based approach and visual tool to support requirements engineering), refactoring of designs using patterns [32], software integration [33], model-based security [34], and performance assessment [35].

Model are useful since humans can review/audit/improve an explicit representation of their systems. But as models grow in complexity, it becomes difficult for a manual analysis to reveal all their subtleties. Hence, many researchers propose support environments to help explore the increasingly complex models that engineers are developing. Gray, et al, [36] have developed the Constraint-Specification Aspect Weaver (C-Saw) that uses aspect-oriented approaches [37] to help engineers in the process of model transformation. Cai and Sullivan [38] describe a formal method and tool called *Simon* that “supports interactive construction of formal models, derives and displays design structure matrices ... and supports simple design impact analysis.” Other tools of note are lightweight formal methods such as ALLOY [39] and SCR [40] as well as various UML tools that allow for the execution of life cycle specifications (e.g. the CADENA scenario editor [41]).

Recently, AI has been successful applied to model-based SE. For example, Whittle uses deductive learners to generate lower-level UML designs (state charts) from higher-level constructs (use case diagrams) [42]. More generally, the field of *search-based SE* augments model-based SE with *meta-heuristic* techniques, like genetic algorithms, simulated annealing, etc., to explore a model. Such heuristic methods are hardly complete but, as Clarke et.al. [43] remark: “...software engineers face problems which consist, not in finding *the* solution, but rather, in engineering an *acceptable* or *near optimal solution* from a large number of alternatives.” [43].

Search-based SE is most often used to optimizing software testing [44–47] but it has had application in numerous other areas. With Feather, we have used search-based SE for requirements analysis [31]. Other researchers [48, 49] use genetic algorithms to examine ways of modularizing software [43] or developing effort estimators [50–52]. In all, Rela [53] lists 123 publications where search-based methods have been applied to the above applications as well as automatic synthesis of software defect predictors; assisting in component design; developing multiprocessor

schedules; re-engineering old systems into a better one; and searching for compiler optimizations.

To use a search-based approach, software engineers have to reformulate their problem by:

- Finding a *representation of the problem* that can be symbolically manipulated (e.g. simulated or mutated). Such representations always exist with model-based SE.
- Defining a *fitness function* (a.k.a. “utility function” or “objective function”); i.e. an “oracle” that scores a model configuration. Current model-based SE methods rarely offer such a function (exception: formal methods that generate temporal constraints). In our experience, generating such a fitness function is usually possible, albeit requiring days of work with the domain experts [54].
- Determining an appropriate set of *manipulation operators* to select future searches based on the prior searches [55].

Data mining is one way to implement automatic manipulation operators. A data miner searches through the space of possible concepts for a combination of concepts that describes some target theory [56]. Given, say, the output from a Monte Carlo simulation of a model, a 21<sup>st</sup> century data miner can sift through gigabytes of data looking for the core concepts that most select for preferred output.

### 3. “Hard Modeling” Problems

Before describing our research into data mining for automatically generating manipulation operators, we must motivate why we don’t use traditional methods. As argued in the section, we seek constraints to model inputs that are *satisficing* (so, not necessary optimal), stable in the face of uncertainty, automatically generated, that reduce cognitive overload, and which work for large non-linear and noisy models. Such constraints are called “solutions to hard modeling problems”.

Traditional methods may not yield such solutions. For example, one simple traditional method is *binary plotting*. After sampling the model, all model inputs might be plotted against the output score in two-dimensional binary plots. The best constraints on the inputs are those that select for the best output. However, if noise in a models yields plots that are not smooth, then it becomes a time-consuming and subjective task to examine binary plots to detect a region of input parameters that most improves the outputs. Also, studying binary plotting will not find conjunctions of multiple input constraints that contribute to better output. Further, as the dimensionality of model input increases, the cognitive effort required to study the plots can become prohibitively expensive.

Many researchers have developed impressive visual environments for decreasing the cognitive overload associated with exploring a multi-dimensional space. For example, Figure 1 shows a tool developed at IBM that augments a standard three-

8 *Menzies, Kiper, Greenwald, Hu, Raffo, Setamanit*

dimensional display with visual cues relating to higher-dimensional data; e.g., supplementary dimensions are shown bottom right; blue circles around the axle show circular motion information; and a color key, shown bottom left, indicates how colors in the display relate to density information [57].

Such visualizations help analysts explore visual information, but they present their own challenges. There are still limits on how many dimensions can be displayed (e.g. we have yet to see effective visualizations for more than a ten dimensionality space). Also, note how the tool shown in Figure 1 contains numerous controls that allow analysts to change various visualization options. As the visualization environment grows more sophisticated, some users find they have traded a data browsing problem with the new problem of exploring the full range of the effects of all the controls.

When manual exploring of options fail, automatic methods can be applied. Optimization packages can be applied to data or the equations of a system to find “sweet spots” that maximize the score resulting from model outputs. Related methods include *sensitivity analysis* [1] and *design of experiments* (DOE) methods [58]. A canonical sensitivity analysis method might be to compute eigenvectors of a linear system in order to understand its long-term temporal behavior [2,3]. As for design of experiments, DOE exercises an existing model and helps shed light on the response surface of the model. DOE does identify gradients and key parameters for a model.

While useful for some models, these automatic methods do not apply to all models. For example, optimization methods can fail for non-linear models. Any model with an “if” statement introduces a “cliff” where the effects of inputs on outputs abruptly changes. For such models, there is no linear continuous solution

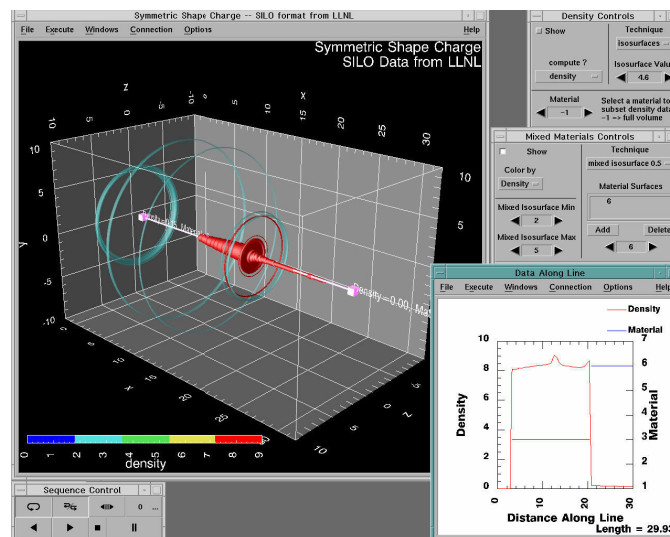


Fig. 1. A visualization tool for scientific data. From [57].



that applies either side of the “cliff”. Model uncertainty also complicates sensitivity analysis. For example, the eigenvector technique described would yield spurious results if the coefficients on the models are not known with certainty. Lastly, a DOE analysis can be complicated by the dimensionality, noise, and visualization problems described above.

Our own view on hard modeling is that there exists a “loophole” in search problems that can make hard problems far easier to manage. We call this loophole “collars and clumps”.

#### 4. Collars and Clumps

This research assumes that many models can be controlled by a small number of key variables which we call *collars*. Collars restrict the behavior of a model such that their state space *clumps*; i.e. only small number of states are used at runtime. If so, the output of a data miner could be simplified to constrain just the collar variables that switch the system between a few clumps.

To visualize collars, imagine an execution trace spreading out across a program. Initially, the trace is very small and includes only the inputs. Later, the trace spreads wider as *upstream* variables effect more of the *downstream* variables (and the inputs are the most *upstream* variables of all). At some time after the appearance of the inputs, the variables hold a set of values. Some of these values were derived from the inputs while others may be default settings that, as yet, have not been affected by the inputs. The union of those values at time  $t$  is called the *state*  $s_t$  of the program at that time.

Multiple execution traces are generated when the program is run multiple times with different inputs. These traces reach different branches of the program. Those branches are selected by tests on conditionals at the root of each branch. The *controllers* of a program are the variables that have different settings at the roots of different branches in different traces. Programs have *collars* when a handful of the controllers in an early state  $s_t$  control the settings of the majority of the variables seen in later states .

A related effect to *collars* is *clumping*. If a program has  $v$  variables with range  $r$ , then the maximum number of states is  $r^v$ . Programs *clump* when most of those states are never *used* at runtime; i.e.  $|used|/(r^v) \approx 0$ . Clumps can cause collars:

- The size of *used* is the cardinality cross product of the ranges seen in the controllers.
- If that cardinality is large, many states will be generated and programs won't clump.
- But if that cross product is small, then the deltas between the states will be small – in which case controlling a small number of collar variables would suffice for selecting what states are reached at runtime.


There is much theoretical and empirical evidence for expecting that many models

10 *Menzies, Kiper, Greenwald, Hu, Raffo, Setamanit*

often contain *collars* and *clumps*.

#### 4.1. Theoretical Evidence

With Singh [59], we have shown that collars are an expected property of Monte Carlo simulations where the output has been discretized into a small number of output classes. After such a discretization, many of the inputs would reach the same goal state, albeit by different routes. The following diagram shows two possible distinct execution paths within a Monte Carlo simulation both leading to the same goal; i.e.  $a \rightarrow goal$  or  $b \rightarrow goal$ .



Each of the terms in lower case in the above diagram represent a probability of some event; i.e.  $0 \leq \{a_i, b_i, c, d, goal_i\} \leq 1$ . For the two pathways to reach the *goal*, they must satisfy the collar  $M$  or the larger collar  $N$  (each collar is a conjunction). As the size of  $N$  grows, the product  $\prod_{j=1}^N b_j$  decreases and it becomes less likely that a random Monte Carlo simulation will take steps of the larger collar  $N$ .

The magnitude of this effect is quite remarkable. Under a variety of conditions, the narrower collar is thousands to millions of times more likely. For example, when  $|M| = 2$  and  $N > M$ , the condition for selecting the larger collar is  $\frac{d}{c} \geq 64$ ; i.e. the larger collar  $N$  will be used only when the  $d$  pathway is dozens of times more likely than  $c$ . The effect is more pronounced as  $|M|$  grows; at  $|M| = 3$  and  $N > M$ , the condition is  $\frac{d}{c} \geq 1728$ ; i.e. to select the larger collar  $N$ , the  $d$  pathway must be thousands of times more likely than  $c$  (for more details, see [59]). That is, when the output space is discretized into a small number of outputs, and there are multiple ways to get to the same output, then a randomized simulation (e.g. a Monte Carlo simulation) will naturally select for small collars.

While the mathematics may be arcane, the intuition is simple. Suppose all the power goes out in a street of hotels. If all those hotels were designed by different architects then their internal search spaces would be different (number of rooms per floor, distance from each room to a flight of stairs, number of stairwells, etc). After half an hour, some of the guests bumble around in the dark and get outside to the street. Amongst the guests that have reached the street, there would be more guests from hotels with simpler internal search spaces (fewer rooms per floor, less distance from each room to the stairs, more stairwells).

As to *clumping*, Druzdel [60] observed this effect in a medical monitoring system. The system had 525,312 possible internal states. However, at runtime, very few were

ever reached. In fact, the system remained in one state 52% of the time, and a mere 49 states were used, 91% percent of the time. Druzdel showed mathematically that there is nothing unusual about his application. If a model has  $n$  variables, each with its own assignment probability distribution of  $p_i$ , then the probability that the model will fall into a particular state is  $p = p_1 p_2 p_3 \dots p_n = \prod_{i=1}^n p_i$ . By taking logs of both sides, this equation becomes

$$\ln p = \ln \prod_{i=1}^n p_i = \sum_{i=1}^n \ln p_i \quad (1)$$

The asymptotic behavior of such a sum of random variables is addressed by the central limit theorem. In the case where we know very little about a model,  $p_i$  is uniform and many states are possible. However, the *more* we know about the model, the *less* likely it is that the distributions are uniform. Given enough variance in the individual priors and conditional probabilities or  $p_i$ , the expected case is that the frequency with which we reach states will exhibit a log-normal distribution; i.e. a small fraction of states can be expected to cover a large portion of the total probability space; and the remaining states have practically negligible probability.

The assertion that many types of models display this clumping behavior is quite important for the style of data mining (treatment learning) that we advocate. In application to a clumping model with collars, Monte Carlo simulation, followed by TAR2, suffices to summarize that model in an effective way:

- TAR2's rules never need to be bigger than the collars. Hence, if the collars are small, TAR2's rules can also be small.
- If a model clumps, then, very quickly, a Monte Carlo simulation would sample most of the reachable states. TAR2's summarization of that simulation would then include most of the important details of a model.

#### 4.2. Empirical Evidence

Empirical evidence for clumps first appeared in the 1950s. Writing in 1959, Samuel studied machine learning for the game of checkers [61]. At the heart of his program was a 32-term polynomial that scored different configurations. For example, *king center control* means that a king occupies one of the center positions. The program learned weights for these variable coefficients. After 42 games, the program had learned that 12 variables were important, although only 5 of these were of any real significance.

Decades later, we can assert that deleting irrelevant variable has proven to be a useful strategy in many domains. For example, Kohavi and John report experiments on 8 real world datasets where, on average, 81% of the non-collar variables can be ignored without degrading the performance of a model automatically learned from the data [62].

If models contain collars, or if the internal state space clumps, then much of the reachable parts of a program can be reached very quickly. This *early coverage*

effect has been observed many times. In a telecommunications application, Avritzer, Ros, & Weyuker found that a sample of 6% of all inputs to this system covered 99% of all inputs seen in about one year of operation (and a sample of just over 12% covered 99.9%) [63]. Further evidence for early coverage can be found in the mutation testing literature. In mutation testing, some part of a program is replaced with a syntactically valid, but randomly selected, variant (e.g. switching “less than” signs to “greater than”). This method of testing is useful for getting an estimate of what percentage of errors have been discovered by testing. Wong compared results using  $X\%$  of a library of mutators, randomly selected ( $X \in \{10, 15, \dots, 40, 100\}$ ). Most of what could be learned from the program could be learned using only  $X=10\%$  of the mutators; i.e. after a very small number of mutators, new mutators acted in the same manner as previously used mutators [64]. The same observation has been made elsewhere by Budd [65] and Acree [66].

If the space of possible execution pathways within a program are limited, then program execution would be observed to clump since it could only ever repeat a few simple patterns. Empirically such limitations have been observed in procedural and declarative systems. Bieman and Schultz [67] report that 10 or fewer paths through programs explored 83% of the du-pathways (a du-path is a set of statements in a computer program from a definition to a use of a variable. This is one common form of structural coverage testing.) Harrold [68] studied the control graphs of 4000 Fortran routines and 3147 C functions. Potentially, the size of a control graph may grow as the square of the number of statements (in the case where every statement was linked to every other statement.) This research found that, in these case studies, the size of the control graph is a linear function of the number of statements. In an analogous result, Pelánek reviewed the structures of dozens of formal models and concluded that the internal structure of those models was remarkably simple: “state spaces are usually sparse, without hubs, with one large SCC [strongly connected component], with small diameter<sup>b</sup> and small SCC quotient”<sup>c</sup> [69]. This sparseness of state spaces was observed previously by Holtzmann where he estimate the average degree of a vertex in a state space to be 2 [70].

Pelánek hypothesises that these “observed properties of state spaces are not the result of the way state spaces are generated nor of some features of specification languages but rather of the way humans design/model systems” [69]. Pelánek does not expand on this, but we assert that generally SE models are simple enough to be controlled by treatment learning since they were written by humans with limited short-term memories [71] who have difficulty designing overly-complex models.

<sup>b</sup>The diameter of a graph (of a state space here) is the number of edges on the largest shortest path between any two vertices.

<sup>c</sup>SCC quotient is a measure of the complexity of a graph.

## 5. Data Mining with Collars and Clumps

The TAR2 *treatment learner* [72, 73] is a data miner that is specialized for generating models containing only collar variables. TAR2 finds the difference between classes. Formally, the algorithm is a *contrast set learner* [74, 75] that uses *weighted classes* [76] to steer the inference towards the preferred behavior. We call TAR2's output "treatments" since the minimal rules generated by the algorithm are similar to medical treatment policies that try to achieve the most benefit, with the least intervention. The core intuition of TAR2 is that it is unnecessary to search for the collars— they will reveal themselves after some limited random sampling. To see that, recall that collar variables control the settings in the rest of the system. Any execution trace that reaches a goal must pass through the collars (by definition). Therefore, to find the collars, all an algorithm needs to do is find the attribute ranges with very different frequencies in traces that reach different goals.

Detecting collars via this sampling method is very simple to implement. Consider a log of golf playing behavior shown in Figure 2. This log contains four attributes (outlook, temperature, humidity, wind) and 3 classes (none, some, lots) that convey the amount of golf played. We recommend an exponential scoring system for the classes, starting at two<sup>d</sup>. For example, our golfer could weight the classes in Figure 2 as *none*=2 (worst), *some*=4, *lots*=8 (best).

TAR2 seeks attribute ranges that occur frequently in the highly weighted classes and rare in the lower weighted classes. Let *a.r* be some attribute range e.g. *outlook.overcast* means that the outlook is for overcast skies.  $\Delta_{a.r}$  is a heuristic measure of the worth of *a.r* to improve the frequency of the *best* class.  $\Delta_{a.r}$  uses the following definitions:

*X(a.r)*: is the number of occurrences of that attribute range in class *X*; e.g. in this data *lots(outlook.sunny)*=2 since there are 2 cases with outlook = *sunny* and class = *lots*.

*all(a.r)*: is the total number of occurrences of that attribute range in all classes; e.g. *all(outlook.sunny)*=5.

*best*: the highest scoring class; e.g. *best* = *lots*;

*rest*: the non-best class; e.g. *rest* = {*none*, *some*};

*weight*: The weight of a class *X* is symbolized by  $\$X$ ; (Thus,  $\$best = 8$ .)

$\Delta_{a.r}$  is calculated as follows:

$$\Delta_{a.r} = \frac{\sum_{X \in rest} (\$best - \$X) * (best(a.r) - X(a.r))}{all(a.r)}$$

<sup>d</sup>If the weights run, say, {bad=0,ok=1,good=2} then the difference from *bad* to *ok* scores the same as *ok* to *good*. An exponentially weighting scheme, starting at two, finds greater and greater rewards moving to better classes. For further details, see [77].

14 *Menzies, Kiper, Greenwald, Hu, Raffo, Setamanit*

When  $a.r$  is *outlook.overcast*, then  $\Delta_{outlook.overcast}$  is calculated as follows:

$$\frac{\overbrace{((8-2) * (4-0))}^{lots \rightarrow none} + \overbrace{((8-4) * (4-0))}^{lots \rightarrow some}}{4 + 0 + 0} = \frac{40}{4} = 10$$

<i>outlook</i>	<i>temp</i> (°F)	<i>humidity</i>	<i>windy?</i>	<i>class</i>	<i>weight</i>
<i>sunny</i>	85	86	<i>false</i>	<i>none</i>	2
<i>sunny</i>	80	90	<i>true</i>	<i>none</i>	2
<i>sunny</i>	72	95	<i>false</i>	<i>none</i>	2
<i>rain</i>	65	70	<i>true</i>	<i>none</i>	2
<i>rain</i>	71	96	<i>true</i>	<i>none</i>	2
<i>rain</i>	70	96	<i>false</i>	<i>some</i>	4
<i>rain</i>	68	80	<i>false</i>	<i>some</i>	4
<i>rain</i>	75	80	<i>false</i>	<i>some</i>	4
<i>sunny</i>	69	70	<i>false</i>	<i>lots</i>	8
<i>sunny</i>	75	70	<i>true</i>	<i>lots</i>	8
<i>overcast</i>	83	88	<i>false</i>	<i>lots</i>	8
<i>overcast</i>	64	65	<i>true</i>	<i>lots</i>	8
<i>overcast</i>	72	90	<i>true</i>	<i>lots</i>	8
<i>overcast</i>	81	75	<i>false</i>	<i>lots</i>	8

Fig. 2. A log of some golf-playing behavior.

To *build* a treatment, TAR2 explores combinations of attribute ranges up to some user-specified maximum size  $s$  (where the size  $s$  is the number of attribute ranges in a conjunction of attributes). Given  $n$  attributes, the size of this search is  $\frac{n!}{s!(n-s)!}$ . To make this search feasible, TAR2 must keep  $s$  small. Therefore, TAR2 first assesses each attribute range, in isolation, i.e., with  $s = 1$ . A preliminary pass builds one singleton treatment for each attribute range. The attribute ranges are then scored by the  $\Delta$  of these singleton treatments. Treatment generation is constrained to just the attribute ranges with a score greater than a user-supplied threshold (default value= 1; maximum useful value yet found= 7).




Fig. 3. Finding treatments that can improve golf playing behavior. With no treatments, we only play lots of golf in  $\frac{6}{5+3+6} = 57\%$  of cases. However, assuming *outlook=overcast*, we play golf lots of times in 100% of cases.

To *apply* a treatment, TAR2 rejects all example entries that contradict the conjunction of the attribute ranges in the treatment. E.g., if the treatment was

domain	# rows	# columns		#class	time(sec)
		numeric #	discrete #		
iris	150	4	0	3	< 1
wine	178	13	0	3	< 1
car	1,728	0	6	4	< 1
autompg	398	6	1	4	1
housing	506	13	0	4	1
pageblocks	5,473	10	0	5	2
cocomo	30,000	0	23	4	2
reachness	25,000	4	9	4	3
circuit	35,228	0	18	10	4
reachness2	250,000	4	9	4	23
pilot	30,000	0	99	9	86

Fig. 4. Runtimes for TAR2 on different domains. First 6 data sets come from the UC Irvine machine learning data repository [78]; “cocomo” comes from a COCOMO software cost estimation model [79]; “pilot” comes from the NASA Jet Propulsion Laboratory [80].

$humidity \geq 85 \wedge windy = true$ , then 11 of the lines of Figure 2 would be rejected. The ratio of classes in the remaining examples is compared to the ratio of classes in the original example set (in the humidity and wind treatment just given, this ratio would be 3/14). The *best treatment* is the one that most increases the relative percentage of preferred classes. In our golf example, a single best treatment was generated containing *outlook=overcast*. Figure 3 shows the class distribution before and after that treatment. That is, if we select a vacation location with *overcast* weather, then we should be playing *lots* of golf, all the time.

In practice, despite the  $\frac{n!}{s!(n-s)!}$  search, TAR2 scales well. Figure 4 shows TAR2’s runtime on 11 data sets with varying numbers of rows and columns. Running on a relatively slow machine (a 333 MHz Windows machine with 512MB of ram), TAR2 terminated in tens of seconds, even on data sets with up to 250,000 rows, each with nearly 100 attributes.

## 6. Case Studies

In theory, we expect that many models contain collars and clumps. If so, tools like TAR2 should be able to find tiny treatments that control the behavior of the models. This section tests that theory on several case studies.

### 6.1. Inspection Policies

The first case study contrasts treatment learning with traditional learners. It will be seen that treatments are dramatically smaller, and more understandable, than the model learned by standard data miners.

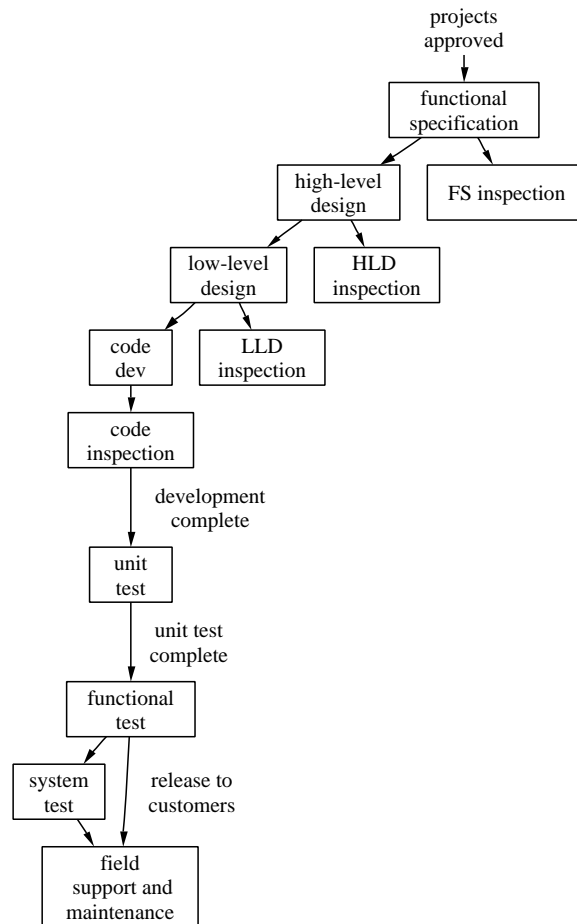


Fig. 5. High-level block diagram of a discrete event model of one company's software process.

Figure 5 offers a high-level view of a quantitative software process model [81]. At each phase of that process, inspections are conducted of the functional specification (FS), high-level design (HLD), low-level design (LLD), and the code (CODE). Raffo modeled these phases, and the inspections using a Statemate<sup>TM</sup> state-based simulation model and an Extend<sup>TM</sup> discrete event model containing 30+ process steps with two levels of hierarchy. Some of the inputs to the simulation model included productivity rates for various processes; the volume of work (i.e. KSLOC); defect detection and injection rates for all phases; effort allocation percentages across all phases of the project; rework costs across all phases; parameters for process overlap; the amount and effect of training provided; and resource constraints.

Model outputs are the development *expense* (person months), product *quality*



(number of high severity defects) and project *duration* (calendar months) which are combined as follows:

$$utility = 40 * (14 - quality) + 320 * (70 - expense) + 640 * (24 - duration) \quad (2)$$

The justification for this style of utility function is discussed in detail in [81]. In summary, this function was created after extensive debriefing of the business users.

The model has been extensively validated. The model's process diagrams, model inputs, model parameters and outputs were reviewed by members of the software engineering process group as well as senior developers and managers. In other studies, the model was used to accurately predict the performance of several past releases of the project. Finally, in *special case* studies, the model was used to predict unanticipated special cases. Specifically, when predicting the impact of developing overly complex functionality, the model predicted that development would take approximately double the normal development schedule. Initially rejected by management, it was later found that this model's predictions corresponded quite accurately with the company's actual experience.

In this example, we will use the model to assess different software inspection policies. For each phase of Figure 5, four types of inspections can potentially be applied. These four types are listed below, sorted by their cost and effectiveness. For example, *full Fagan inspections* are most expensive and find the most issues. At the other end of the scale, doing *no inspections* is cheapest but finds no issues:

- F: A *full Fagan inspection* [82] is a seven step process with pre-determined roles for inspection participants. For the company studied by Raffo, the *defect detection capability*<sup>e</sup> of their full Fagan inspections was  $TR(0.35, 0.50, 0.65)$ <sup>f</sup>. Such studies use between 4 and 6 staff, plus the author of the artifact being inspected.
- B: A *baseline inspection* is a continuation of current practice at the company under study. The baseline inspection at this company was essentially a poorly performed Fagan inspection, Historical records show that these baseline inspections have varying defect detection capabilities of  $\{min, mode, max\} = \{0.13, 0.21, 0.30\}$ .
- W: *Walk through* inspections conducted informally by an outside consultant. Historical records show that these inspections have a defect detection capability of  $TR(0.07, 0.15, 0.23)$ .
- N: *No inspection*;

Each type of inspection can be performed at each phase; i.e. there are  $4^4$  possible inspection polices. A data set for TAR2 was prepared by running each configuration 50 times; i.e.  $50 * 4^4 = 12800$  samples. Each run was then tagged with its utility,

<sup>e</sup>Defect detection capability is the percentage of defects that are latent in the artifact that is being inspected that are detected.

<sup>f</sup> $TR(a, b, c)$  denotes a triangular distribution with minimum, mode, max of  $a, b, c$  respectively.




Fig. 6. Utility frequency in 12800 samples, *before* and *after* treatment.

using Equation 2. These utilities were discretized into four classes, of approximately equal frequency:

- *class1* : value of Equation 2 < 9843
- *class2* :  $9843 \leq$  value of Equation 2 < 10698
- *class3* :  $10698 \leq$  value of Equation 2 < 11664
- *class4* :  $11664 \leq$  value of Equation 2  $\leq$  14755

The *before* bar chart in Figure 6 shows the frequency of these classes in the untreated model. Note that many (40%) of the samples generate the lowest *class1* utility.

TAR2 was then applied to learn treatments that distinguish the desired *class4* utilities from the rest. The best treatment generated by the algorithm was very small only recommended changing one of the inspection processes; i.e.

$$hidesign_{12} = F$$

The *after* bar chart of Figure 6 shows the effect of imposing this treatment of  $hidesign_{12}=F^{\S}$  (full Fagan inspections for high-level designs) onto the simulator. The treated model is much improved: it generates zero *class1* and *class2* outputs and, in 80% of cases, generates *class4* outputs. Further, the improvement was achieved without having to control the inspection policies in the other phases.

If we were not concerned with finding minimal solutions, TAR2 would have been called again on data generated from the inspection model, after the inputs have been constrained to  $hidesign_{12}=F$ . This iterative process could repeat until it was shown that further constraints did not improve the output. Such *interactive*

<sup>\S</sup>In this data set, each attribute is labeled with its column number so *hidesign<sub>12</sub>* appears in the twelfth column the input.

C4.5 and CART are *iterative dichotomization* learners that seek the best attribute value *splitter* that most simplifies the data that fall into the different splits. Each such splitter becomes a root of a tree. Sub-trees are generated by calling iterative dichotomization recursively on each of the splits.

CART is defined for continuous target concepts and its *splitters* strive to reduce the standard deviation of the data that falls into each split. C4.5 is defined for discrete class classification and uses an information-theoretic measure to describe the diversity of classes within a data set.

A leaf generated by CART stores the average value of the class selected by the branch while a leaf generated by C4.5 stores the most frequency class. Hence, C4.5 is called a *decision tree* learner while CART is called a *regression tree*.

Fig. 7. About C4.5 and CART

*treatment learning* has been applied on other models (e.g. see [80]). However, for the sake of exposition, the execution of this model is not explored further.

In terms of the discussion in §3, an important feature of *hidesign.12=F* is that it is stable despite uncertainty in other parameters. Despite large scale variation of all other parameters in this model, this treatment yielded the effect seen in Figure 6. In terms of supporting commercial practices, this is a very useful result. Large corporations may have little impact on their satellite organizations or contractors. Hence, they often have to carefully select what policies to implement across the company. In terms of Equation 2, the treatment learned in this example representing the *least action* that offers the *most reward*.

Also, in terms of advocating treatment learning, the most important feature of this example is what is *missing*. To learn its treatments, TAR2 imported samples with 51 variables (50 inputs and one utility score). It then generated treatments, the best of which, used only one of the inputs. Other data miners, such as C4.5 and CART (described in Figure 7), are far less succinct. For example, Figure 8 show the results of passing the inspection model data to C4.5 [83]. Figure 9 and Figure 10 pass shows the output from the CART and linear regression learners:

- C4.5 and CART use the *iterative dichotimization* algorithm described in Figure 7.
- Linear regression tries to fit one straight line through the observed values. The line offers a set of predicted values and the distance from these predicted values to the actual values is a measure of the error associated with that line. Linear regression tools such as the least squares regression package search for a line that minimizes that sum of the squares of the error.
- C4.5 predicts for discrete class symbols (e.g. *class1*, *class2*, *class3*, *class4*) so this algorithm used the same data as TAR2.
- Linear regression and CART make numeric predictions. These algorithms used the TAR2 data with the *classN* symbols replaced by raw numerics of Equation 2.

A comparison of the output from TAR2, C4.5 (in Figure 8), and results from CART and linear regression (in Figure 9 and Figure 10) demonstrate that standard methods of summarizing data (linear regression, decision trees, regression trees) can generate much larger theories than treatment learning. The reason for this is very simple. Theories learned from iterative dichotomization describe the features that separate all of the target variables. However, treatments just describe the minimal deltas *between* preferred and undesirable targets.

Another advantage of treatment learning is that it is much easier to derive ac-

```

corErr_gg50 <= 72: class1
corErr_gg50 > 72
|   corErr_gg50 <= 154
|   |   spec_gg2 = b
|   |   |   hiDist_gg13 = 0
|   |   |   |   corErr_gg11 <= 17: class1
|   |   |   |   corErr_gg11 > 17: class2
|   |   |   |   hiDist_gg13 = tr
|   |   |   |   |   corErr_gg31 <= 18: class1
|   |   |   |   |   corErr_gg31 > 18
|   |   |   |   |   |   inspEff_gg19 <= 33.742185: class1
|   |   |   |   |   |   inspEff_gg19 > 33.742185: class2
|   |   |   |   spec_gg2 = f
|   |   |   |   |   corErr_gg21 <= 22
|   |   |   |   |   |   detCap_gg38 <= 0.563188
|   |   |   |   |   |   |   inspDur_gg10 <= 117.043523
|   |   |   |   |   |   |   |   corErr_gg50 <= 120: class1
|   |   |   |   |   |   |   |   corErr_gg50 > 120
|   |   |   |   |   |   |   |   |   inspDur_gg40 <= 292.445892: class2
|   |   |   |   |   |   |   |   |   inspDur_gg40 > 292.445892: class1
|   |   |   |   |   |   |   |   |   inspDur_gg10 > 117.043523: class1
|   |   |   |   |   |   |   |   |   detCap_gg38 > 0.563188: class2
|   |   |   |   |   |   |   corErr_gg21 > 22: class2
|   |   |   |   spec_gg2 = n
|   |   |   |   |   inspEff_gg29 <= 336.924853: class2
|   |   |   |   |   inspEff_gg29 > 336.924853: class1
|   |   |   |   spec_gg2 = w
|   |   |   |   |   corErr_gg11 <= 3: class1
|   |   |   |   |   corErr_gg11 > 3
|   |   |   |   |   |   corErr_gg11 <= 4
|   |   |   |   |   |   |   inspEff_gg9 <= 32.423212: class2
|   |   |   |   |   |   |   inspEff_gg9 > 32.423212
|   |   |   |   |   |   |   |   inspDur_gg10 <= 35.939846
|   |   |   |   |   |   |   |   |   detCap_gg8 <= 0.110608: class2
|   |   |   |   |   |   |   |   |   detCap_gg8 > 0.110608: class1
|   |   |   |   |   |   |   |   |   inspDur_gg10 > 35.939846: class2
|   |   |   |   |   |   |   corErr_gg11 > 4
|   |   |   |   |   |   |   |   detCap_gg8 <= 0.145107
|   |   |   |   |   |   |   |   |   inspEff_gg9 <= 49.227938: class1
|   |   |   |   |   |   |   |   |   inspEff_gg9 > 49.227938
|   |   |   |   |   |   |   |   |   |   detCap_gg8 <= 0.131955: class2
|   |   |   |   |   |   |   |   |   |   detCap_gg8 > 0.131955: class1
|   |   |   |   |   |   |   |   |   |   detCap_gg8 > 0.145107
|   |   |   |   |   |   |   |   |   |   |   inspEff_gg9 <= 33.741217: class1
|   |   |   |   |   |   |   |   |   |   |   inspEff_gg9 > 33.741217: class2
|   |   |   |   corErr_gg50 > 154
|   |   |   |   |   detCap_gg8 <= 0.618237: class2
|   |   |   |   |   |   detCap_gg8 > 0.618237
|   |   |   |   |   |   |   inspDur_gg30 <= 137.500433: class3
|   |   |   |   |   |   |   |   inspDur_gg30 > 137.500433: class2

```

Fig. 8. A decision tree learned via C4.5.

```

policy=NNFN,BBNF,WNNF,NNWF,BBBB,FFFF,NNNF,BFFF <= 0.5 :
| policy=NNWN,NNNW,WNNW,BNNW,WNNB,BNNB,WWW,NNNB,NNNW, FNNF,NNWB,NNFN,BBNF,WNNF,NNWF,BBBB,FFFF,NNNF,BFFF <= 0.5 :
| | policy=WNNN,NWNN,FNNW,FNNB,BNNN,NNWN,NNNW,WNNW,BNNW,WNNB,BNNB,WWW,NNNB,NNNW,FNNF,NNWB,NNFN,BBNF,WNNF,
  NNWF,BBBB,FFFF,NNNF,BFFF <= 0.5 :
| | | _detCap_8 <= 0.431 : 6940
| | | _detCap_8 > 0.431 : 8180
| | policy=WNNN,NWNN,FNNW,FNNB,BNNN,NNWN,NNNW,WNNW,BNNW,WNNB,BNNB,WWW,NNNB,NNNW,FNNF,NNWB,NNFN,BBNF,WNNF,
  NNWF,BBBB,FFFF,NNNF,BFFF > 0.5 :
| | | _detCap_8 <= 0.491 :
| | | | _inspDur_10 <= 94 : 8780
| | | | _inspDur_10 > 94 : 7860
| | | _detCap_8 > 0.491 : 9320
| | policy=NNWN,NNNW,WNNW,BNNW,WNNB,BNNB,WWW,NNNB,NNNW,FNNF,NNWB,NNFN,BBNF,WNNF,NNWF,BBBB,FFFF,NNNF,BFFF > 0.5 :
| | | policy=BNNB,WWW,NNNB,NNNW,FNNF,NNWB,NNFN,BBNF,WNNF,NNWF,BBBB,FFFF,NNNF,BFFF <= 0.5 : 9750
| | | policy=BNNB,WWW,NNNB,NNNW,FNNF,NNWB,NNFN,BBNF,WNNF,NNWF,BBBB,FFFF,NNNF,BFFF > 0.5 : 10300
policy=NNFN,BBNF,WNNF,NNWF,BBBB,FFFF,NNNF,BFFF > 0.5 :
| policy=NNWF,BBBB,FFFF,NNNF,BFFF <= 0.5 :
| | _detCap_38 <= 0.418 : 10700
| | _detCap_38 > 0.418 : 11400
| | policy=NNWF,BBBB,FFFF,NNNF,BFFF > 0.5 :
| | | _corErr_50 <= 155 : 11900
| | | _corErr_50 > 155 :
| | | | policy=NNNF,BFFF <= 0.5 : 12400
| | | | policy=NNNF,BFFF > 0.5 :
| | | | _detCap_38 <= 0.519 : 12500
| | | | _detCap_38 > 0.519 : 13600

```

Fig. 9. A regression tree learned by CART from data generated from the Figure 5 model. In this figure, and Figure 10, policies are denoted as a four-part sequence describing the inspection policy at each phase. For example, “FFFF” denotes using no inspections for code but full Fagan inspections for all earlier phases.

```

utility = 7370
+ 1220policy=WNNN,NWNN,FNNW,FNNB,BNNN,NNWN,NNNW,WNNW,BNNW,WNNB,BNNB,WWW,
  NNNB,NNNW,FNNF,NNWB,NNFN,BBNF,WNNF,NNWF,BBBB,FFFF,NNNF,BFFF
+ 258policy=FNNB,BNNN,NNWN,NNNW,WNNW,BNNW,WNNB,BNNB,WWW,NNNB,NNNW, FNNF,
  NNWB,NNFN,BBNF,WNNF,NNWF,BBBB,FFFF,NNNF,BFFF
+ 640policy=NNWN,NNNW,WNNW,BNNW,WNNB,BNNB,WWW,NNNB,NNNW,FNNF,NNWB,NNFN,BBNF,WNNF,NNWF,BBBB,FFFF,NNNF,BFFF
+ 505policy=NNNW,WNNW,BNNW,WNNB,BNNB,WWW,NNNB,NNNW,FNNF,NNWB,NNFN,BBNF,WNNF,NNWF,BBBB,FFFF,NNNF,BFFF
- 190policy=WNNW,BNNW,WNNB,BNNB,WWW,NNNB,NNNW,FNNF,NNWB,NNFN,BBNF,WNNF,NNWF,BBBB,FFFF,NNNF,BFFF
+ 206policy=BNNB,WWW,NNNB,NNNW,FNNF,NNWB,NNFN,BBNF,WNNF,NNWF,BBBB,FFFF,NNNF,BFFF
- 177policy=WWW,NNNB,NNNW,FNNF,NNWB,NNFN,BBNF,WNNF,NNWF,BBBB,FFFF,NNNF,BFFF
+ 890policy=NNNB,NNNW,FNNF,NNWB,NNFN,BBNF,WNNF,NNWF,BBBB,FFFF,NNNF,BFFF
- 296policy=NNNW,FNNF,NNWB,NNFN,BBNF,WNNF,NNWF,BBBB,FFFF,NNNF,BFFF
+ 545policy=NNFN,BBNF,WNNF,NNWF,BBBB,FFFF,NNNF,BFFF
+ 640policy=WNNF,NNWF,BBBB,FFFF,NNNF,BFFF
+ 459policy=NNWF,BBBB,FFFF,NNNF,BFFF - 143policy=BBBB,FFFF,NNNF,BFFF
+ 835policy=NNNF,BFFF - 142_spec_2=n,b + 27_spec_2=b + 3690_detCap_8
- 135_inspEff_9 - 7320_inspDur_10 - 64400_corErr_11 - 760_detCap_18
- 134_inspEff_19 - 7310_inspDur_20 - 64300_corErr_21 - 9100_detCap_28
- 134_inspEff_29 - 7320_inspDur_30 - 64300_corErr_31 + 6170_detCap_38
- 134_inspEff_39 - 7320_inspDur_40 - 64400_corErr_41
+ 21500_inspEff_48 + 1.17e6_inspDur_49 + 64400_corErr_50

```

Fig. 10. Linear regression learned from data generated from the Figure 5 model.

tions from treatments than from the standard methods shown in Figures 8, 10, 9. To be sure, decision trees can be analyzed to find branch values that most selected for preferred classes while most discarding undesired classes (the initial TAR2 prototype was such a post-processor). However, TAR2 achieves the same result directly without the need to interface to another learner.

## 6.2. Studying the CMM

The previous study explored a numeric model where all the influences were precisely specified. This second case study takes a numeric model and adds a large degree of uncertainty in the numerics. This second study shows that, even in presence of large degrees of uncertainty, TAR2 can still find useful treatments.

An important feature of this second study is that it analyzes a class of models that can defeat standard methods. The model contains dozens of if-then rules; i.e. it is neither linear nor continuous: small changes in the environment can lead to “cliffs” where the model behavior changes abruptly. Also, the model contains non-deterministic choices (see the *rany* operator, discussed below) and so its behavior can be highly noisy.

This study uses a rule-based model of the costs and benefits model of CMM level 2 (hereafter, CMM2) [84, p125-191]. We elected to study CMM2 since, in our experience, many organizations can achieve at least this level. CMM2 is less concerned with issues of (e.g.) which design pattern to apply, than with what overall project structure should be implemented. Improving CMM2-style decisions is important since in early software life cycle, many CMM2-style decisions affect the resource allocation for the rest of the project.

CMM2 was encoded using the JANE propositional rule-based language [85]. JANE’s rules take the form *Goal if SubGoals* such as the one shown in Figure 11.

```

stableRequirements
  if   effectiveReviews
  and  requirementsUsed
  and  sEteamParticipatesInPlanning
  and  documentedRequirements
  and  sQAactivities
  and  (reviewRequirementChanges
        rany softwareConfigurationManagement
        rany baselineChangesControlled
        rany workProductsIdentified
        rany softwareTracking
      ).

```

Fig. 11. Part of CMM2, encoded in the JANE language.

JANE is a backward chaining language: to prove a *Goal*, JANE tries to find rules that prove each of the *SubGoals*. Each *SubGoal* contributes some *Cost* and *Chances* to the *Goal*. JANE’s *Chances* define the extent to which a belief in one vertex can propagate to another. *Costs* let an analyst model the common situation where some of the *Cost* of some procedure is amortized by reusing its results many times. Hence, the *first* time we use a proposition, we incur its *Cost* but afterwards, that proposition is free of charge.

The *Cost* and *Chances* of a proposition are either provided by the JANE programmer or computed at runtime via a traversal of the rules:

- When searching *X if not A*, the *Chances* of *X* are  $1 - \text{Chances}(A)$  and

$$\text{Cost}(X) = \text{Cost}(A).$$

- When searching  $X$  if  $A$  and  $B$  and  $C$ , the *Chances* and *Costs* of  $X$  are (respectively) the product of the chances and the sum of the costs of  $A, B, C$ .
- When searching  $X$  if  $A$  or  $B$  or  $C$ , then the *Cost* and *Chances* of  $X$  are taken from the first member of  $A, B, C$  that is satisfied.

These *and*, *or*, *not* operators can be insufficient to capture the decision making of business users. For examples, in our experience, business users select CMM2 options, often in a somewhat arbitrary manner. To model this, JANE includes a *rany* operator (short for “random any”):

- The *rany* operator is like *or* except that (e.g.)  $X$  if  $A$  *rany*  $B$  *rany*  $C$  succeeds if some random number of  $A, B, C$  (greater than one) succeeds. Unlike *and*, *or* which explore their operands in a left-to-right order, *rany* explores its *SubGoals* in a random order. If at least one succeeds, then the *Cost* and *Chances* of  $X$  is the sum and product (respectively) of the *Cost* and *Chances* of the satisfied members of  $A, B, C$ .

*Rany* is useful when searching for subsets that contribute to some conclusion. For example, the JANE rule in Figure 11 offers several essential features of *stableRequirements* plus several optional factors relating to monitoring change in evolving projects – the essential features are *and*-ed together while the optional factors are *rany*-ed together.

Figure 11 includes 11 propositions. Our model of CMM2, written in JANE, has 55 propositions ( $range = \{t, f\}$ ). Of those 55 propositions, 27 were identified by

<i>baselineAudits, baselineChangesControlled, changeRequestsHandled, changesCommunicated, configurationItemStatusRecorded, deviationsDocumented, documentedDevelopmentPlan, documentedProjectPlan, earlyPlanning, formalReviewsAtMilestones, goodUnitTesting, identifiedWorkProducts, periodicSoftwareReviews</i>	<i>planRevised, requirementsReview, requirementsUsed, reviewRequirementChanges, risksTracked, SCMplan, SCMplanUsed, SElifeCycleDefined, SEteamParticipatesInPlanning, SEteamParticipatesOnProposal, SQAauditsProducts, SQAplan, SQAplanUsed, SQAreviewActivities, workProductsIdentified</i>
---	--

Fig. 12. Management actions in the CMM2 model. SQA= software quality assurance and SCM= software configuration management)


24 *Menzies, Kiper, Greenwald, Hu, Raffo, Setamanit*

Fig. 13. Ratios of different software project types seen in four situations.

our users as actions that could be changed by managers (see Figure 12).

Apart from *rany*, JANE supports one other mechanism for exploring the space of possibilities within CMM2. When defining *Costs* and *Chances*, the programmer can supply a *range* and a *skew*. For example:

```
goodUnitTesting and cost = 1 to +5
```

defines the *cost* of *goodUnitTesting* as being somewhere in the range 1 to 5, with the mean skewed slightly towards 5 (denoted by the “+”).

Similarly, while all the *Chances* values were based on expert judgment, their precise value is subjective. Hence, each such *Chances* value  $X$  was altered to be a range

```
chances = 0.7*X to 1.3*X
```

During a simulation, the *first* time a *Cost* or *Chance* is accessed, it is assigned randomly according to the range and skew. The assignment is cached so that all subsequent accesses use the same randomly generated value. After each simulation, the cache is cleared. After thousands of simulations, JANE can sample the “what-if” behavior resulting from different assignments within the range and many different *rany* choices.

Data from 2000 simulations was passed from the CMM2 model to TAR2. Each simulation was classified into one of four classes:

- *class=0*: High cost, low chance;
- *class=1*: Low cost, low chance;
- *class=2*: High cost, high chance;



$T_1$ : *requirementsUsed.Cost=lower and  
 not periodicSoftware-Reviews and  
 formalReviewsAtMilestones.Cost=lower*  
 $T_2$ : *requirementsUsed.Cost=lower and  
 goodUnitTesting.Cost=middle and  
 formalReviewsAtMilestones.Cost=lower*  
 $T_3$ : *goodUnitTesting.Cost=lower and  
 periodicSoftwareReviews.Cost=middle and  
 formalReviewsAtMilestones.Cost=lower*

Fig. 14. The three best treatments found in the CMM2 model.

- *class=3*: Low cost, high chance.

That is, our preferred projects are cheap and highly likely while expensive, low odds projects are to be avoided.

Figure 13 shows three sets of actions learned by TAR2. The right-hand-side histogram shows the baseline distributions seen in the 2000 simulations. The other histograms show how those ratios change after applying the treatments learned by TAR2; The *worth* of each option is a reflection of the proportion of good and bad projects, compared to the baseline, i.e. ( $worth(baseline) = 1$ ). Note that as *worth* increases, the proportion of preferred projects also increases.

Figure 14 shows the three best treatments ( $T_1, T_2, T_3$ ) found using this technique (and Figure 13 compared the effects of these treatments to the untreated examples). Note that the values of each attribute are reported using the tags *no*, *lower*, *middle*, or *upper*. In treatment learning, continuous attribute ranges are divided into N-discrete bands based on percentile positions. For N=3, we can name the bands *lower*, *middle*, *upper* for the lower, middle, and upper 33% percentile bands.

In Figure 14, the treatments are advising to lower the cost of:

- *Using requirements*: This could be accomplished by (e.g.) sharing them around the development team in some search-able hypertext format
- *Performing formal reviews at milestones*: This could be accomplished by (e.g.) using ultra-lightweight formal methods such as proposed by Leveson [86].
- *Performing good unit testing*: This could be accomplished by (e.g.) hiring better test engineers.

An interesting feature of Figure 14 is what is *missing*:

- None of the treatments proposed adjusting the *Chances* of any action. In this study, changing *Cost* will suffice.
- Of the 27 actions listed in in Figure 12, only the four underlined actions appear in the top three treatments. That is, management commitment to undertake  $27-4=23$  of the actions is less use-

ful than changing on *formalReviewsAtMilestones*, *goodUnitTesting*, *periodicSoftwareReviews*, and *requirementsUsed*

- The value *not* in  $T_1$  is a recommendation against *periodicSoftwareReviews* (plus lowering the costs of using requirements and formal reviews at milestones). Note that if *periodicSoftwareReviews* are conducted,  $T_3$  is saying that there is no apparent need to reduce the cost of such reviews.

More generally, in a result consistent with the prior studies, despite the uncertainties introduced by *rany* and the *cost/chances* ranges, TAR2 found a small number of CMM2 process options that have a significant impact on the project.

Note that the conclusions of Figure 14 are not general to all software projects. The *Chances* values used in this study came from some local domain knowledge about the likelihood that process change *A* will effect process change *B*. The *Cost* values were domain-specific as well. In other organizations, with different work practices and staff, those *Chances* and *Cost* values could be very different.

### 6.3. Other Case Studies

Treatment learning has been applied to spacecraft design to find how to cover more requirements, reduce risk, at the least cost [80, 87]. It has also been applied to software process control using:

- A Chung-Mypolopous soft-goal graph to find better coverage of the non-functional requirements [88].
- COCOMO effort and risk models models to find options selecting for lower effort and fewer risks [79];
- COCOMO effort, risk, and defect prediction models models to find project options selecting for lower effort and fewer threats and lower defects [89];
- Qualitative inference diagrams to find requirements selecting for higher quality [90].
- The NASA SILAP model (that selects V&V tasks in order to most lower risks) [91];

Treatment learning has also been applied to:

- Finite state machines to find topologies that reduce the CPU cost of applying formal methods [92, 93].
- Models of the global economy so study methods of extending human life expectancy [94];
- Maximizing whiskey production [95];

In all the case studies explored by TAR2, the same three observations were made:

- Treatment learning can find very small treatments, even for seemingly complex models;

- These treatments can be far smaller than models generated by standard data miners.
- Despite uncertainties or variabilities in the model, TAR2 was able to find effective treatments that selected for preferred model output (but the less uncertainty or model variability, the smaller the variance in TAR2's predicted output for the treated model).

## 7. Conclusion

Understanding model configuration options means understanding how input choices affect output scores. That understanding is complex for a certain class of *hard models*; i.e. those with high dimensionality models that are non-linear, non-continuous and built in domains with much noise or other uncertainties, and where managers have limited control over all model inputs.

Hard modeling problems may defeat standard methods. Visualization can't handle very large dimensionality. Analytical methods such as an eigenvector study offer spurious results if the parameters of the variables are uncertain. Other standard automatic methods may be defeated by "cliffs" in non-linear models where the association between inputs and outputs changes abruptly. Data mining methods can handle non-linear models and scale to very large dimensionality. Sadly, data mining techniques like neural nets, genetic algorithms, C4.5 and CART can yield models that are incomprehensible to humans.

TAR2 is a special kind of data miner that produces very succinct output. It assumes that within models there exists a small number of key variables that control the rest. There is much evidence for this assumption. The mathematics of clumps and collars promises that models naturally contain structures that greatly restrict the space of possible model behaviors. TAR2 is a data miner designed to exploit such collars and clumps. It is a minimal contrast set learner that returns a "treatment"; i.e. a minimal, most influential set of deltas between different classes of outcome. The case studies in this paper show that a minimal list of the differences between concepts can be *much smaller* than a detailed description of all aspects of a concept. For models where TAR2 can generate succinct summaries, its algorithm can significantly improve searched-based methods of data mining.

TAR2 addresses the hard modeling problems (discussed in §3) as follows:

- TAR2 offers minimal constraints on the input space and tracks the effects of those constraints, while letting all the other variables vary randomly. Hence, its proposed solutions are not brittle to changes outside the treatments.
- Since it only references a subset of the model inputs, it is a dimensionality reduction tool. In this report, we offered examples where TAR2 reasoned over 100-variable inputs spaces. Elsewhere, we have run it on data sets with over 250 variables. In all cases seen to date, it reduces those spaces to a handful of variables.
- Such small solutions are easier to explain and audit than solutions using all

model inputs.

- Further, when managers do not have the budget or authority to control all model input variables, TAR2 can offer them a minimal set which they can use to focus their resources.
- TAR2 has been applied to models with large amounts of noise. For example, in a TAR2-style analysis, we often explore what happens to the solutions when the variance on the model variables increases. Such studies return statements of the form “the solutions offered by this analysis hold for variances up to the following critical threshold values, after which we do not know how to control this domain”.

Treatment learning is not indicated for low-dimensionality linear continuous models built in domains that have no noise or other uncertainties, and where managers have full control over all model inputs. Other reasons not to use our tools include when there is no need to explain or audit models, where reducing model variance is not valuable, and there exists budgets for building and using maximal models.

As to further work, there is no reason to polarize the SE modeling field into “traditional methods” vs “treatment learning”. Much could be achieved by combining the two techniques. For example, many of the methods described in §3 suffered from the curse of dimensionality. TAR2 could be used as a fast dimensionality reduction tool that could focus a data visualization environment or a sensitivity analysis on the parts of the inputs space that are most crucial. Ideally, that focusing need not wait till the simulation terminates. In *incremental treatment learning*, TAR2 offers feedback during a simulation run in order to guide the simulator into regions of interest. Before TAR2 can be deployed in that manner, it must be optimized so that it can run fast enough to keep up with the simulator. Currently, we are exploring stochastic methods for that optimization.

## References

1. Saltelli, A., Chan, K., and Scott, E. *Sensitivity Analysis*. Wiley, (2000).
2. Ishida, Y. In *Proceedings of IJCAI '89*, 1174–1179., (1989).
3. Levins, R. and Puccia, C. *Qualitative Modeling of Complex Systems: An Introduction to Loop Analysis and Time Averaging*. Harvard University Press, Cambridge, Mass., (1985).
4. Mintzberg, H. *Harvard Business Review*, 29–61 July-August (1975).
5. Simon, H. A. *Models of bounded rationality*, volume 2. MIT Press, (1982).
6. Rosenbloom, P., Laird, J., and Newell, A. *The SOAR Papers*. The MIT Press, (1993).
7. Miller, A. *Subset Selection in Regression (second edition)*. Chapman & Hall, (2002).
8. Seidewitz, E. *IEEE Software* **20**(5), 26–32 Sept.-Oct. (2003).
9. Sendall, S. and Kozacaynski, W. *IEEE Software* **20**(5), 42–45 Sept.-Oct. (2003).
10. b. hailpern and p. tarr. *ibm systems journal* **45**(3), 451–461 (2006).
11. Object Management Group. *MDA Guide Version 1.0.1*, June (2003).
12. Greenfield, J. and Short, K. *Software factories : assembling applications with patterns, models, frameworks, and tools*. Wiley Publishing, Indianapolis, IN, (2004).

13. Waddington, D. and Lardieri, P. *IEEE Computer* **39**(2), 28–29 February. (2006).
14. Clancey, W., Sachs, P., Sierhuis, M., and van Hoof, R. In *Proceedings PKAW '96: Pacific Knowledge Acquisition Workshop*, Compton, P., Mizoguchi, R., Motoda, H., and Menzies, T., editors. Department of Artificial Intelligence, (1996).
15. Law, A. and Kelton, B. *Simulation Modeling and Analysis*. McGraw Hill, (2000).
16. Harrell, H., Ghosh, L., and Bowden, S. *Simulation Using ProModel*. McGraw-Hill, (2000).
17. Kelton, D., Sadowski, R., and Sadowski, D. *Simulation with Arena, second edition*. McGraw-Hill, (2002).
18. Raffo, D. and Menzies, T. *Journal of Information, Software and Technology* **47**(15), 1009–1017 December (2005).
19. Raffo, D. and Menzies, T. In *Proceedings of the 6th International Workshop on Software Process Simulation Modeling (ProSim'05)*, (2005).
20. Abdel-Hamid, T. and Madnick, S. *Software Project Dynamics: An Integrated Approach*. Prentice-Hall Software Series, (1991).
21. Serman, H. *Business Dynamics: Systems Thinking and Modeling for a Complex World*. Irwin McGraw-Hill, (2000).
22. Akhavi, M. and Wilson, W. In *Proceedings of the 5th Software Engineering Process Group National Meeting (Held at Costa Mesa, California, April 26 - 29)*. Software engineering Institute, Carnegie Mellon University, (1993).
23. Harel, D. *IEEE Transactions on Software Engineering* **16**(4), 403–414 April (1990).
24. Martin, R. and Raffo, D. M. *International Journal of Software Process Improvement and Practice* June/July (2000).
25. Martin, R. and D. M, R. *Journal of Systems and Software* **59**(3) (2001).
26. Donzelli, P. and Iazeolla, G. *Journal of Systems and Software* **59**(3) December (2001).
27. Setamanit, S., Wakeland, W., and D.Raffo. *Software Process: Improvement and Practice, (Forthcoming)* (2007).
28. Bratko, I. *Prolog Programming for Artificial Intelligence. (third edition)*. Addison-Wesley, (2001).
29. Iwasaki, Y. In *The Handbook of Artificial Intelligence*, A. Barr, P. C. and Feigenbaum, E., editors, volume 4, 323–413. Addison Wesley (1989).
30. Mi, P. and Scacchi, W. *IEEE Transactions on Knowledge and Data Engineering* , 283–294 September (1990).
31. Feather, M. and Cornfordi, S. *Requirements Engineering Journal* **8**(4), 248–265 (2003).
32. France, R., Ghosh, S., Song, E., and Kim, D. *IEEE Software* **20**(5), 52–58 Sept.-Oct. (2003).
33. Denno, P., Steves, M. P., Libes, D., and Barkmeyer, E. J. *IEEE Software* **20**(5), 59–63 Sept.-Oct. (2003).
34. Jerjens, J. and Fox, J. In *ICSE '06: Proceeding of the 28th international conference on Software engineering*, 819–822 (ACM Press, New York, NY, USA, 2006).
35. Balsamo, S., Marco, A. D., Inverardi, P., and Simeoni, M. *IEEE Transactions on Software Engineering* **30**(5) May (2004).
36. Gray, J., Lin, Y., and Zhang, J. *IEEE Computer* **39**(2), 51–58 February (2006).
37. Filman, R. E. *Aspect-Oriented Software Development*. Addison-Wesley, Boston, (2004).
38. Cai, Y. and Sullivan, K. J. In *ASE '05: Proceedings of the 20th IEEE/ACM international Conference on Automated software engineering*, 329–332 (ACM Press, New York, NY, USA, 2005).
39. Jackson, D. *ACM Trans. Softw. Eng. Methodol.* **11**(2), 256–290 (2002).
40. Heitmeyer, C. In *Encyclo-*

30 *Menzies, Kiper, Greenwald, Hu, Raffo, Setamanit*

*pedia of Software Engineering*, Marciniak, J. J., editor, January (2002). Available from <http://chacs.nrl.navy.mil/publications/CHACS/2002/2002heitmeyer-encse.pdf>

41. Childs, A., Greenwald, J., Jung, G., Hoosier, M., and Hatcliff, J. *IEEE Computer* **39**(2) February (2006). Available from <http://projects.cis.ksu.edu/docman/view.php/7/129/CALM-Cadena-IEEE-Computer-Feb-2006.pdf>
42. Whittle, J. and Jayaraman, P. In *IEEE International Conference on Requirements Engineering (RE2006)*, (2006).
43. Clarke, J., Dolado, J., Harman, M., Hierons, R., Jones, B., Lumkin, M., Mitchell, B., Mancoridis, S., Rees, K., Roper, M., and Shepperd, M. *IEE Proceedings-Software* **150**(3), 161–175 (2003).
44. Jones, B., Sthamer, H.-H., and Eyres, D. *Software Engineering Journal* **11**, 299–306 (1996).
45. Jones, B., Eyres, D., and Sthamer, H.-H. *Computer Journal* **41**(2), 98–107 (1998).
46. Pargas, R., Harrold, M., and Peck, R. R. *Journal of Software Testing, Verification and Reliability* **9**, 263–282 (1999).
47. Tracey, N., Clarke, J., and Mander, K. In *International Symposium on Software Testing and Analysis*, 73–81. *ACM/SIGSOFT*, March (1998).
48. Harman, M., Hierons, R., and Proctor, M. In *GECO 2002: Proceedings of the Genetic and Evolutionary Computation Conference*, 1351–1358. *Morgan Kaufmann*, July (2002).
49. Lutz, R. *Journal of Systems Architecture* **47**, 613–634 (2001).
50. Aguilar-Ruiz, J., Ramos, I., Riquelme, J., and Toro, M. *Information and Software Technology* **43**(14), 875–882 December (2001).
51. Dolado, J. J. *IEEE Transactions of Software Engineering* **26**(10), 1006–1021 (2000).
52. Dolado, J. J. *Information and Software Technology* **43**, 61–72 (2001).
53. Rela, L. *Master's thesis, Lappeenranta University of Technology*, (2004).
54. Menzies, T. *International Journal of Human-Computer Studies*, special issue on evaluation of KE techniques **51**(4), 783–799 October (1999). Available from <http://menzies.us/pdf/99csm.pdf>
55. Harman, M. and Jones, B. *Journal of Information and Software Technology* **43**, 833–839 December (2001).
56. Mitchell, T. *Machine Learning*. McGraw-Hill, (1997).
57. Treinish, L. (1998). *IBM Research Center, Yorktown Heights, NY*. Available from [http://www.research.ibm.com/people/l/lloyd/dm/function/dm\\_fn.htm](http://www.research.ibm.com/people/l/lloyd/dm/function/dm_fn.htm)
58. Boning, D. and Mozumder, P. *IEEE Transactions on Semiconductor Manufacturing* **7**(2), 233–244 May (1994).
59. Menzies, T. and Singh, H. In *Soft Computing in Software Engineering*, Madravio, M., editor. Springer-Verlag, (2003). Available from <http://menzies.us/pdf/03maybe.pdf>
60. Druzdzal, M. In *Proceedings of the Tenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-94)*, 187–194, (1994). Available from <http://www.pitt.edu/AFShome/d/r/druzdzal/public/html/abstracts/uai94.html>
61. Samuel, A. L. *IBM Journal* **3**(3), 211–229 July (1959).
62. Kohavi, R. and John, G. H. *Artificial Intelligence* **97**(1-2), 273–324 (1997).

63. Avritzer, A., Ros, J., and Weyuker, E. IEEE Software , 76–82 September (1996).
64. Wong, W. and Mathur, A. The Journal of Systems and Software **31**(3), 185–196 December (1995).
65. Budd, T. Mutation analysis of programs test data. *PhD thesis, Yale University, (1980).*
66. Acree, A. On Mutations. *PhD thesis, School of Information and Computer Science, Georgia Institute of Technology, (1980).*
67. Bieman, J. and Schultz, J. Software Engineering Journal **7**(1), 43–51 (1992).
68. Harrold, M., Jones, J., and Rothermel, G. Empirical Software Engineering **3**, 203–211 (1998).
69. Pelanek, R. In Proceedings SPIN'04 Workshop, (2004). Available from [http://www.fi.muni.cz/~xpelanek/publications/state\\_spaces.ps](http://www.fi.muni.cz/~xpelanek/publications/state_spaces.ps)
70. Holzmann, G. J. ATT Technical Journal **69**(2), 32–44 (1990).
71. Müller, G. The Psychological Review **63**, 81–97 (1956). Available from <http://www.well.com/~smalin/miller.html>
72. Menzies, T. IEEE Intelligent Systems (2003). Available from <http://menzies.us/pdf/03aipride.pdf>
73. Menzies, T. and Hu, Y. In Artificial Intelligence Review, (2007). Available from <http://menzies.us/pdf/07tar2.pdf>
74. Bay, S. and Pazzani, M. In Proceedings of the Fifth International Conference on Knowledge Discovery and Data Mining, (1999). Available from <http://www.ics.uci.edu/~pazzani/Publications/stucco.pdf>
75. Webb, G. I., Butler, S., and Newlands, D. In KDD '03: Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining, 256–265 (ACM Press, New York, NY, USA, 2003).
76. Cai, C., Fu, A., Cheng, C., and Kwong, W. In Proceedings of International Database Engineering and Applications Symposium (IDEAS 98), August (1998). Available from [http://www.cse.cuhk.edu.hk/kdd/assoc\\_rule/paper.pdf](http://www.cse.cuhk.edu.hk/kdd/assoc_rule/paper.pdf)
77. Hu, Y. Master's thesis, Department of Electrical Engineering, University of British Columbia, (2003). Masters Thesis.
78. Blake, C. and Merz, C. (1998). URL: <http://www.ics.uci.edu/~mlearn/MLRepository.html>
79. Menzies, T. and Sinsel, E. In Proceedings ASE 2000, (2000). Available from <http://menzies.us/pdf/00ase.pdf>
80. Feather, M. and Menzies, T. In IEEE Joint Conference On Requirements Engineering ICRE'02 and RE'02, 9-13th September, University of Essen, Germany, (2002). Available from <http://menzies.us/pdf/02re02.pdf>
81. Raffo, D. May (1996). *Ph.D. thesis, Manufacturing and Operations Systems.*
82. Fagan, M. IEEE Trans. on Software Engineering , 744–751 July (1986).
83. Quinlan, R. Machine Learning **1**, 81–106 (1986).
84. Paulk, M., Weber, C., Curtis, B., and Chriss, M. The Capability Maturity Model:

32 *Menzies, Kiper, Greenwald, Hu, Raffo, Setamanit*

Guidelines for Improving the Software Process. *Addison-Wesley*, (1995).

85. *Menzies, T. and Kiper, J. In ASE-2001, (2001). Available from <http://menzies.us/pdf/01ase.pdf>*
86. *Leveson, N., Cha, S., and Shimall, T. IEEE Software 8(7), 48-59 July (1991).*
87. *Cornford, S. L., Feather, M. S., Dunphy, J., Salcedo, J., and Menzies, T. In Proceedings of the IEEE Aerospace Conference, Big Sky, Montana, (2003). Available from <http://menzies.us/pdf/03aero.pdf>*
88. *Chiang, E. and Menzies, T. Software Process: Improvement and Practice 7(3-4), 141-159 (2003). Available from <http://menzies.us/pdf/03spip.pdf>*
89. *Menzies, T. and Richardson, J. In COCOMO forum, 2005, (2005). Available from [http://menzies.us/pdf/05xomo\\_cocomo\\_forum.pdf](http://menzies.us/pdf/05xomo_cocomo_forum.pdf)  
class="hC".*
90. *Menzies, T. and Richardson, J. IEEE Computer October (2006). Available from <http://menzies.us/pdf/06qrre.pdf>*
91. *Fisher, M. and Menzies, T. In HICSS'06, (2006). Available from <http://menzies.us/pdf/06hicss.pdf>*
92. *Menzies, T., Owen, D., and Cukic, B. In Formal Aspects of Agent-Based Systems, (2002). Available from <http://menzies.us/pdf/02trust.pdf>*
93. *Owen, D., Menzies, T., and Cukic, B. In IEEE Conference on Automated Software Engineering (ASE '02), (2002). Available from <http://menzies.us/pdf/02moretest.pdf>*
94. *Geletko, D. and Menzies, T. In IEEE NASE SEW 2003, (2003). Available from <http://menzies.us/pdf/03radar.pdf>*
95. *Burkleaux, T., Menzies, T., and Owen, D. In Proceedings of WITSE 2005, (2004). Available from <http://menzies.us/pdf/04lean.pdf>*