MyPYTHIA: A Recommendation Portal
for Scientific Software and Services[*]

E. HOUSTIS, A. C. CATLIN, N. DHANJANI and J. R. RICE
Department of Computer Sciences
Purdue University
West Lafayette, IN 47906

N. RAMAKRISHNAN
Department of Computer Science
Virginia Tech University
Blacksburg, VA 24061

V. VERYKIOS
College of Information Science and Technology
Drexel University
Philadelphia, PA 19104

Abstract

In this paper, we outline the design of a recommendation system (MyPYTHIA) implemented as a web portal. MyPYTHIA's design objectives include evaluating the quality and performance of scientific software on grid platforms, creating knowledge about which software and computational services should be selected for solving particular problems, selecting parameters of software (or of computational services) based on user-specified computational objectives, providing access to performance data and knowledge bases over the web, and enabling recommendations for targeted application domains. MyPYTHIA uses a combination of statistical analysis, pattern extraction techniques, and a database of software performance to map feature-based representations of problem instances to appropriate software. MyPYTHIA's open architecture allows the user to customize it for conducting individual case studies. We describe the architecture as well as several scientific domains of knowledge enabled by such case studies.

1

Contents

1 Introduction

The new economic realities require the rapid prototyping of manufactured artifacts and rapid solutions to problems with numerous interrelated elements. Thus, computational modeling and simulation has already become the main tool of 'big' and 'small' science, and it is shifting from the single physical component design to the design of a whole physical system. Increasingly, such simulations are being dynamically composed by scientists working collaboratively on the grid, even though they are geographically distributed [6, 7, 9, 10, 11, 12].

The realization of this simulation scenario requires that grid users be able to identify relevant software and its parameters before implementation on a remote computing resource. However, the obstacles to selecting the best solution approach for a particular problem, and subsequently finding an appropriate software implementation, are often difficult, and sometimes even impossible, to surmount. What is required is a knowledge-based technology for the location and assembly of scientific software components. We refer to such services as 'recommendation services' and the facilities by which they are invoked as recommendation systems or 'recommender systems' [14, 27, 29, 38, 34, 36, 35, 37].

In this paper, we describe the MyPYTHIA web recommendation portal that evolved from our earlier efforts in scientific recommendation systems [14, 38]. The precursor to this research is the original PYTHIA system described in [38]. This facility used exemplar-based reasoning and emphasized the design of learning techniques to make recommendations of PDE methods/software. In [14], we extended this methodology into the more general PYTHIA-II database architecture, with well-defined interfaces for learning and reasoning. The recommendation problem was 'configurable,' in that the knowledge engineer had to define appropriate database records and specify interactions with the underlying execution environment. PYTHIA-II then provided all the necessary enabling technologies to prototype recommender systems. MyPYTHIA is the final step in this evolution and is meant to be a truly web-accessible system that can be customized by any user desiring recommendation services. MyPYTHIA addresses the following design objectives:

- Evaluate the quality and performance of software/machines pairs; MyPYTHIA allows users to incorporate native performance data and provides all necessary facilities for their evaluation and data mining.

- Create knowledge about which software and computational services should be selected for solving particular problems or realizing application components; MyPYTHIA provides data mining procedures to reduce the original data and generate rules from the original performance data.

- Select the parameters of software (or computational service) based on user-specified computational objectives; MyPYTHIA estimates parameters of software/hardware that satisfy given user-defined objectives such as time and/or memory requirements.

- Provide performance data and knowledge for specific software domains over the web; MyPYTHIA can serve as a knowledge base for a priori defined case studies.

- Provide recommendation services for targeted application domains.

Outline of Paper

Section 2 motivates and describes the algorithmic basis of the MyPYTHIA recommendation methodology. Section 3 identifies the major components of MyPYTHIA, along with system implementation details. Section 4 describes three case studies that have been conducted with MyPYTHIA. Section 5 elaborates on one of these studies, providing details of the analysis and descriptions of the results. Section 6 identifies opportunities for future research.

2 MyPYTHIA Methodology

One of the basic research issues in designing recommendation services is understanding the fundamental processes by which knowledge about scientific problems and their solutions is created, validated, and communicated. Some of this knowledge will come directly from experts — scientists and engineers — in the field. Other knowledge can be mined from experimental data arising from simulations. Yet further knowledge will be learned from the experience gained by the recommendation service itself as it extracts performance knowledge about software components running on the grid and applied to various problems. It is thus not surprising that the word 'recommendation' is interpreted differently in the scope of different projects.

For instance, a polyalgorithm approach [31, 32], where a decision procedure is hardwired to choose among different solution methods, can be viewed as a first level of recommendation service. It is, however, a very restrictive one since the decision procedure is statically hardwired to the application and, except in artificial problem domains, does not generalize to new situations. In many cases, it does not take into account all available problem features and cannot incorporate the latest solution methods. In addition, the recommendations are often 'over-cautious' to prevent various numerical disasters [1]. Furthermore, recommendations provided in this way cannot be used independently of the accompanying software/modules (and hence do not constitute a grid service by themselves). Performance databases such as the PDS server [22] address the problem of mathematical software evaluation, but their main focus is on profiling solution methods and do not aid in problem solving beyond the tuning of architectural parameters. A richer variety of features are modeled by rating services and online taxonomies of mathematical software, but such services are more information-based and cannot be easily invoked in a computational pipeline due to differences in vocabulary and mismatches in problem domain representation.

Before we describe the MyPYTHIA methodology, it is useful to define some notation. We denote by

- $\mathcal{P}$ a collection of problems or problem space,

- $\mathcal{F}$ a feature space for problems in $\mathcal{P}$,

- $\mathcal{A}$ a collection of parameterized scientific software for simulating a specific class of problems in $\mathcal{P}$,

- $\mathcal{M}$ a collection of parameterized machines (hardware),

- $\mathcal{V}$ a set of n performance metrics,

- s a mapping from $\mathcal{P} \times \mathcal{A} \times \mathcal{M}$ to $\mathcal{U}^n \subset \mathcal{R}^n$ where the points in the range of this mapping denote the measurements for the metrics in $\mathcal{V}$ obtained by running (a, m) pairs $((a, m) \in \mathcal{A} \times \mathcal{M})$ on the benchmark space $\mathcal{P}$, and

- $\mathcal{W}$ a collection of profiles for each software/machine pair defined from pairs of metrics from $\mathcal{V}$.

Profiles are high-level summarizations of performance data that are useful in analysis and mining operations. An example of a profile is a low-degree polynomial approximation which can be used to obtain interpolated data over predefined regions. MyPYTHIA's methodology addresses the partition (clustering) of the performance data space $\mathcal{U}^n$ into subclasses $\{\mathcal{C}_i\}_{i=1}^k$ with respect to a set of profiles (or pairs of metrics) and the description (classification) of those classes in terms of problem features from $\mathcal{F}$. The classification part can be used to identify the best (a, m) pair for all $p \in \mathcal{P}$ or a user defined subclass of P. We refer to this problem as the *method/machine selection* problem. In addition, MyPYTHIA can address the problem of determining the associated parameters of the best (a, m) pair (*parameter estimation* problem) along with predicting the performance of the selected pair to the user's problem by a *performance mapping*.

Assuming a 'dense' benchmark of problems from a targeted application domain, MyPYTHIA uses a four-pronged processing phase to realize the solution of the above problems: (i) feature classification of benchmark

problems and applicable software, (ii) performance evaluation of scientific software, (iii) statistical analysis of performance data, and (iv) the (automatic) knowledge discovery of recommendation rules from software performance. These phases have already been defined in previous papers [14, 38, 16] and address many of the difficulties in algorithm recommendation outlined above. In the web portal implementation of MyPYTHIA, only the interface of the portal and the underlying database facilities have been redesigned. The new design allows users to import their own data necessary for the MyPYTHIA methodology to conduct performance evaluation and for the creation of knowledge for a class of software/machine pairs.

3   MyPYTHIA System Components

We now provide detailed descriptions of MyPYTHIA's system components. The MyPYTHIA portal is served by an Apache webserver[1]. PHP[2] is the scripting language of choice and enables a modular design with its support for object oriented class definitions. PHP's built-in support for sessions helps provide the user with an easy to navigate interface. Throughout, we use the term "method" to refer to software that simulates a targeted class of problems.

3.1   Data Management and Retrieval

Each MyPYTHIA user is provided with an isolated view of their own database. Connection to the underlying MySQL (open source[3], RDBMS) database is supported by available PHP database connectivity drivers. Each time new users attempt to view their database contents for the first time, a simple SQL script is run which populates the isolated database with relevant default data. The end-user web-enabled database interface is made possible by tailoring the PhpMyAdmin package[4] to suit the security needs of the MyPYTHIA implementation. The MyPythia database interface indicating the MySQL and PHP supporting components is shown in Figure 1.

The database schema which is used to model and manage the data for the MyPYTHIA recommender has been designed in a very general way, so that the analysis and data mining process can easily be extended to any knowledge domain. Performance data must be generated, collected, and classified outside of the MyPYTHIA system. This data is then imported into the MyPYTHIA storage manager using an XML-like interface. This eliminates the 'domain-specific' representations required to internally generate performance data [14]. Our earlier schema (described in [14]) required users to define all domain-specific elements of the problems and methods so that the system itself could create and execute the programs for performance data acquisition. This assumption is relaxed in MyPYTHIA.

The underlying representations required by MyPYTHIA to support the generation of inference data are listed below.

- *Benchmark Problem Population* :

```
name         varchar() -- problem name (primary key)
description  text      -- high level description of the problem
code         text      -- user description of the problem code
features     text      -- keys for accessing problem characteristics
featurevals  text      -- values for problem characteristics
```

Note that the *code* field is used for descriptive purposes only. Users may enter a description of the code, a path to the code source, or even the code itself in this field, but MyPYTHIA does not use it in the analysis process.

---

[1]http://www.apache.org/

[2]http://www.php.net/

[3]http://www.mysql.com/

[4]http://www.phpwizard.net/projects/phpMyAdmin/

Figure 1: MyPYTHIA database homepage for creating the database, importing performance data, and creating records which identify the benchmark problems, the collection of methods, and the features of interest.

The list of features and their values are the essential structures required in the analysis of the benchmark population.

- *Methods* :

```
name         varchar() -- method name (primary key)
description  text      -- high level description of the method
features     text      -- keys for accessing method characteristics
featurevals  text      -- values for method characteristics
measures     text      -- performance measures to use in method evaluation
```

Features and measures are the critical elements required by MyPYTHIA to analyze the methods and generate rules concerning their application.

- *Features* :

```
name            varchar() -- feature name (primary key)
description     text      -- high level description of the feature
possiblevalues  text       -- values assumed by this feature
```

This is a generic schema for features associated with problems or methods. Features may even be defined to correspond directly to a set of performance data. That is, the feature may occur as a result of problem-method execution, and may not be identified specifically with either a problem or method. An example of *possiblevalues* is 'yes, no' for the feature 'problem uses a uniform grid in all dimensions.'

- *Performance Data* :

```
name         varchar() -- performance record  name (primary key)
description  text      -- high level description of the method
problem      text      -- benchmark problem used to generate this data
method       text      -- method/software used to generate this data
features     text      -- problem-method execution characteristics
featurevals  text      -- values for characteristics
machine      text      -- execution environment, machine characteristics
perfarray    text      -- list of performance measure names
perfvalues   text      -- list of performance values
```

Database a2c - table perf

Field	Type	Attributes	Null	Default	Extra	Action
name	varchar(60)		No			Change Drop Primary Index Unique
description	text		Yes			Change Drop Primary Index Unique
problempointers	text		No			Change Drop Primary Index Unique
methodpointers	text		No			Change Drop Primary Index Unique
featurepointers	text		No			Change Drop Primary Index Unique
perfarray	text		Yes			Change Drop Primary Index Unique
valuesarray	text		No			Change Drop Primary Index Unique

Keyname Unique Field Action
PRIMARY Yes   name Drop
[Documentation]

- Print view
- Browse
- Select
- Insert
- Add new field: [1] [At End of Table ▼] [Go]

- Insert textfiles into table
- View dump (schema) of table
    - ⦿ Structure only   ☐ Add 'drop table'   [Go]
    - ○ Structure and data ☐ Send
                            ☐ Complete inserts
    - ○ CSV data   terminated by [;]

Figure 2: End-user interface for importing the performance data.

The problem and method fields are used to access the features associated with the (problem and method) input used in producing the performance data. They may also be accessed by the end-user to get more comprehensive information about the kinds of problems and methods which are the source for the generated data. The performance measures which are available for analyzing the methods are named in *perfarray*, and their values are listed in *perfvalues*.

- *Feature-Measure Context* :

```
name          varchar() -- context record  name (primary key)
description   text      -- high level description of the context
runs          int       -- number of analyzer runs
problemlist   text      -- list of problems classified by run
methodlist    text      -- list of methods to evaluate
measurelist   text      -- list of performance measures to extract
profileparm   text      -- problem-to-method profile definition
featurelist   text      -- list of features to consider
```

The context record determines the strategy for analyzing the data. The data must be collected and sent to the analyzer in a way that quarantees the formulation of rules that can answer the target recommendation queries. The selected problem-method performance data must be classified according to problem parameter sets (one or more parameterized classes are combined to determine one analyzer run) so that the correct grouping of features will be achieved in the ranking of the methods. Finally, the *profileparm* must identify one particular problem parameter to be used in creating the performance profiles. The measure values are collected for all variations of the 'profile parameter' which are associated with a given problem-parameter-method.

The typical usage of MyPYTHIA assumes that a substantial body of performance data already exists for the user's knowledge domain, and that the user can import this data into MyPYTHIA. The performance data should

Figure 3: End-user interface for browsing the performance database.

represent values that quantify the distinguishing characteristics of the collection of methods (speed of execution, accuracy, scalability, etc.) and include all important measures for evaluating which methods are better than others for a given benchmark problem. The benchmark features represent the interesting and important characteristics of problems that scientists are interested in solving. As far as possible, the collection of methods should be applied to all the benchmark problems and the measures should be collected across all applications of the methods (Statistical pre-processing techniques can help overcome any lack of coverage for the underlying problem population). It remains the task of the domain expert to assemble the benchmark problems and the collection of methods, and to identify the salient, relevant features of both problems and methods. Performance measures should capture the significance of altered problem features on method execution.

The features that characterize the problems, the methods and the execution create the *context* of the inference data that subsequent recommendations are based upon. More specifically, a particular instance of inference data (which the end-user chooses for eliciting a recommendation) is based entirely on the *feature-measure context*. The *performance measures* determine the ranking of methods as applied to the selected benchmark population and the *features* must explain (through the MyPYTHIA rules) why the methods behaved differently.

As an example, scientists may want MyPYTHIA to recommend methods that are scalable. A problem feature, 'size,' could be used to identify methods for which performance deteriorates as problem size increases, but only if the domain expert included performance measures for length of execution time over a range of problem sizes.

3.2 Performance Data Generation and Collection

As described earlier, MyPYTHIA factors out the underlying execution environment for characterizing application/software performance.

The simple schema of MyPYTHIA records and the XML-like import format provide the mechanism for populating the user's MyPYTHIA database with large collections of existing performance data. Users can build retrieval scripts with XML filters to apply to their current data storage environment, extracting the targeted data and inserting the XML definition tags. The result is then imported directly into the MyPYTHIA database manager. A sample XML definition for importing a specific instance of performance data into MyPYTHIA is given below.

Figure 4: End-user interface for accessing the performance data.

```
<complexType name="MyPYTHIARecord">
  <element name="Performance" type="varchar()" />
  <element ref="Description" item="singular pde study 10-2 : 99data -3"/>
  <element ref="perfarray" item="number of iterations, elapsed time,
             max-grid-delta, number of equations, problem size,
             abs max error at nodes, discrete L2 error at nodes,
             discrete L2 residual at midpoints, relative error"/>
  <element ref="perfvalues" item=" 120, 1.73, 1.56E-02, 3969, 38909,
             2.06E-04,2.52E-05,2.29E-01, 0.00328"/>
</complexType>
```

The supporting records for problems, methods and features are defined by users through the MyPYTHIA database manager interface. These supporting records, along with the imported performance data drive the data analysis, rule generation and recommendation processes. Figure 2 shows the user interface for importing performance records into a sample MyPythia database. Figure 3 and Figure 4 show MyPythia browsing and edit interfaces.

3.3   Data Mining Subsystems

Data mining encompasses the process of extracting and filtering performance data for statistical analysis, generating method profiles and ranking them, selecting and filtering data for pattern extraction, and at the end, populating the rule base of the recommender system. Two persistent activities conducted in this subsystem are the statistical analysis and pattern extraction.

Statistical Analysis

The task of the analyzer is to rank a set of methods for a selected sample from the problem population, based on *apriori* specified performance criteria. This ranking (ordering) is a function of the user's computational objectives (error, memory, total execution time). The generation of this ordering is based on the performance curves or profiles such as error versus total elapsed time (on a log-log basis) for each method, machine, and problem triple. These profiles are generated by linear least squares approximations to raw performance data. MyPYTHIA runs the analyzer as a separate process, sending it an input file containing performance data and a set of parameters which specify the mode of operation for the analyzer. In MyPYTHIA, we have borrowed the approach presented in [38], where a non-parametric statistical technique is considered for the ranking of methods with respect to a given set of problem instances. The process for ranking the methods uses multiple comparisons and contrast estimators based on Friedman rank sums [13].

Pattern Extraction Algorithms

A variety of induction algorithms have been incorporated into MyPYTHIA; in this section we outline some of the more important categories of induction algorithms.

Clustering: An area where tremendous progress has been made in inductive learning is clustering — a fundamental procedure in pattern recognition that looks for regularities in training exemplars. By revealing associations between individuals of a population, it can be used to provide a compact representation of the input problem domain. The basic goal of clustering is to obtain a $c-$partition of the input data that exhibits categorically homogeneous subsets, where n is the number of training exemplars and $2 \leq c \leq n$. In MyPYTHIA, by representing problem instances as characteristic vectors, we can use clustering to reveal regularities in the underlying problem population. Recommendations of algorithms can be made by testing for membership of problems in given cluster(s) and choosing the algorithm that best satisfies performance criteria (for other problems in the same cluster).

Neural Networks: More general forms of functional relationships can be modeled by neural networks [17], which can approximate any function to any required level of accuracy (perhaps with exponential increase in complexity). Neural networks use one or more layers of intermediate functional elements to model the dependence of an output signal(s) on given input parameters. The task of learning in this scenario is thus to determine the interconnection strengths and threshold biases (weights) that will result in acceptable approximation. While the general problem of determining weights has been shown to be NP-complete, neural networks have emerged as a valuable tool for pattern recognition and function approximation. In MyPYTHIA, neural networks have been used to associate problem characteristic features with method and machine parameters [18]. Given a training set, MyPYTHIA can functionally model the dependence of problem features on solution properties, and which can lead to the recommendation of a suitable algorithm.

Inductive Logic Programming: While neural networks are attribute-value based techniques, more expressive schemes can be obtained by harnessing the representational power of logic (specifically, first-order logic). In this formalism, facts and measurements are represented as logical facts and data mining corresponds to forming an intensional rule that uses relational representations to model the dependence of a 'goal' predicate on certain input predicates (relations). This 'inductive logic programming' (ILP) approach [5] is one of the most popular in MyPYTHIA and with the aid of software such as PROGOL, it can be used to find patterns such as:

```
method(X,'FFT9') :- laplace(X), constcoeff(X).
```

which indicates that the `FFT9` method is best for problem `X` if `X` is Laplacian and has constant coefficients. Notice the easy comprehensibility of the rule which could be later used in diagnostics and What-if analyses. In addition, such rules can also be recursive, a feature which makes ILP amenable to automated program synthesis and discovery of complex relationships such as protein secondary structure [5]. In addition, ILP allows the incorporation of a priori background knowledge, a necessary pre-requisite for mining in complex structured domains.

ILP systems typically take a database of positive examples, negative examples, and background knowledge and attempt to construct a predicate logic formula (such as `method(X,Y)`) so that all (most) positive examples can be logically derived from the background knowledge and no (few) negative examples can be logically derived.

In the most general setting of first-order logic, relational induction as described above is undecidable. A first restriction to function-free horn clauses results in decidability (this is the form of programs used in the programming language PROLOG). Decidability here is with respect to induction; for deduction, first-order logic is semi-decidable. However, ILP is often prohibitively expensive and the practice in MyPYTHIA is to restrict the hypothesis space to a proper subset of first-oder logic.

Other Propositional Representations: Various attribute–value learning systems like ID3 [25], C4.5 [26], IBL [3], which express the discovered knowledge in attribute-value languages, have also been incorporated into MyPYTHIA. These approaches have the representational ability of propositional logic. Propositional languages are limited and do not allow for representing complex structured objects and relations among objects or their components. The domain (background) knowledge that can be used in the learning process is also of a very restricted form, compared to ILP.

For all of these approaches, an intermediate filter converts required information from MyPYTHIA databases to the format required by each of the data mining packages. Recommendations are then made based on the methodology applicable in the particular learning context. For instance, with an ILP system, we attach to each rule all the positive (true) examples that are covered by the rule. After a rule has fired, we apply a nearest neighbor technique to identify the stored example that most closely matches (closest example) the user's example. After the closest example has been identified, we retrieve the performance data for this example and the recommended best method by the fired rule. The performance profiles for the method and the problem are consulted to provide estimates for the method's parameters and the performance indicators. If it is impossible for the system to satisfy the user's imposed constraints, then it uses the weights selected by the user to determine the best solution that satisfies approximately the constraints with respect to the indicated weights.

3.4   Recommendation and Inference Engine

The recommender system is the end user module of the MyPYTHIA system. It is a form of decision support system and is the only module in MyPYTHIA that is case study dependent as well as domain dependent. We will describe how a recommender system has been generated as an interface for the knowledge generated by the ILP approach in the pattern extraction module.

Each rule selected by the data mining program covers a number of positive and negative examples. The set of positive examples covered for each rule along with the rules, is one part of the input given to the recommender system. The recommender system asks the user to specify the features of the problem he wants to solve. It then uses the `CLIPS` inference engine and checks its rule base to find a rule whose left-hand side satisfies the user selected problem features. Every rule that is found to match the problem features specified by the user is selected and is placed into the *agenda*. Rules are sorted in decreasing order based on their generality (number of examples they cover), and the very first rule in the agenda is fired to determine the best algorithm for the problem the user specifies. Since each rule is associated with a set of positive examples that are covered by this rule, the recommender system goes through the list of positive examples associated with the fired rule and retrieves the example that has the most
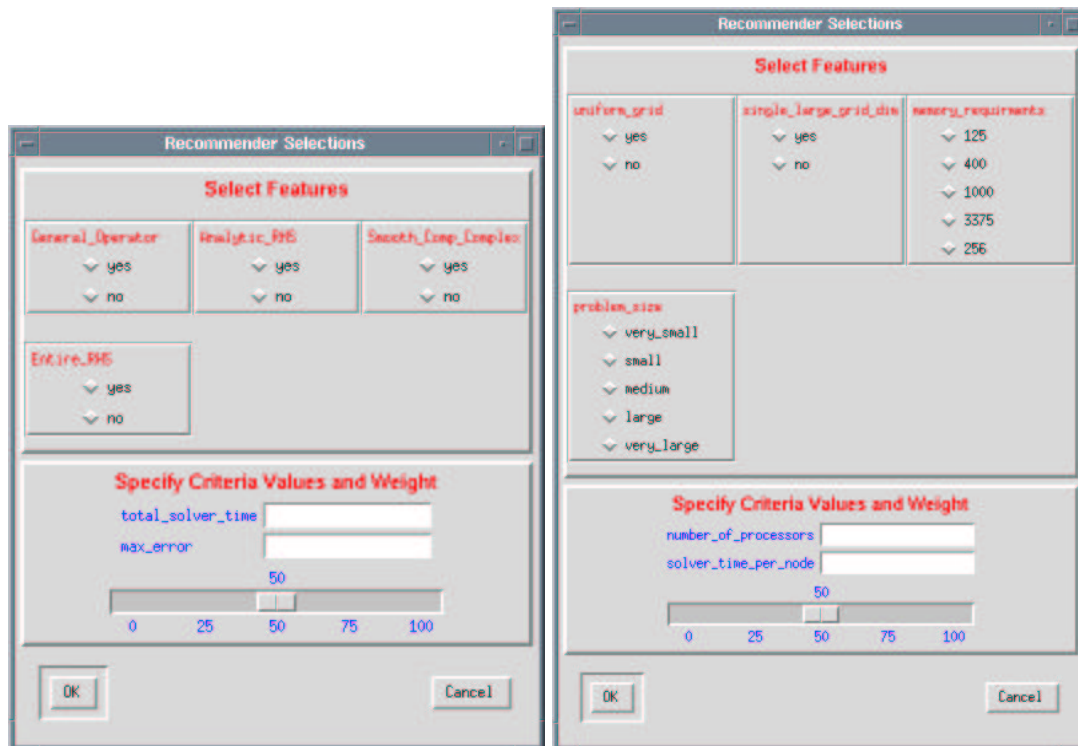
Figure 5: Recommender interfaces for two MyPYTHIA case studies: (left) PDE case study, (right) SWEEP3D application performance characterization.

common features with the user specified problem. This step aids in subsequent parameter estimation.

After this example problem is selected, the fact base of the recommender system is processed in order to provide the user with any required set of parameters for which the user asks advice. The fact base consists of all the raw performance data stored in the database. The recommender system accesses this information by submitting queries generated on the fly, based on the user's objectives and selections. If the user objectives cannot be met, a default recommendation is provided. For some of the case studies presented here, the final step is the recommendation of a certain method as the *best method* to use in satisfing the given conditions. It also indicates what parameter(s) whould be used to achieve the computational objectives imposed by the user.

3.5 User Interface Components of MyPYTHIA

Users access MyPYTHIA through a web-based interface which provides: (i) user management and access authorization, (ii) a web-enabled database interface to the underlying MySQL database, (iii) a guided selection process for identifying how performance data will be classified and retrieved for knowledge creation processing, and (iv) the recommender graphical interface.

When users access MyPYTHIA for the first time (`http://shrine.cs.purdue.edu/pythia3`), they must register for authorization and password assignment. When an authorized, registered user enters the MyPYTHIA environment with id and password, a new database can be created or an already existing database can be selected. The MyPYTHIA database interface supports a full range of functionality for creating, editing, browsing, retrieving and importing MyPYTHIA records. Both graphical and SQL mechanisms are available for managing the stored data.

The specification of the feature-measure context is the starting point for the knowledge building process. MyPYTHIA supports both graphical and SQL methods for identifying which performance data will form the basis for analysis and rules generation. Various graphical selection processes are available, such as

- **Selection by Features:** Users can identify the problem and method features to be studied, and the system will select all performance records associated with problems and methods possessing those features.

- **Selection by Problems and Methods:** Users can select the collection of all problems and methods to be considered in the study. The system will select all performance records naming the chosen problems and methods. Features associated with the selected problems and methods are considered as the basis for the rules generation process.

- **Selection by Performance Data:** Users can select any subset of performance records to consider in the study. Features associated with the problems and methods named in the performance records are considered as the basis for rules generation.

Once the collection of performance records and corresponding features are established, the performance measures are selected. Users select from the available measures listed in the *perfarray* field of the performance record collection. The selected measures must be available for all performance records in the collection, since these measures apply to all analyzer runs. The measures are used to evaluate which methods are better than others, and the measures become the basis for the computational objectives which users specify when requesting software recommendations. The final phase of the feature-measure context specification is the grouping of performance records for analyzer runs. The information specified for the feature-measure context is saved into the context record, and is then used to guide the data retrieval for analysis and method ranking. Users can also browse the context record to realize the basic input to MyPYTHIA's inference data generation process.

The recommendation interface is created dynamically, using the inference context to identify the relevant features and measures. Examples of the MyPYTHIA recommender interface for two application case studies are shown in Figure 5.

4 Current Applications Coverage of MyPYTHIA

We now describe the chief application case studies that have used MyPYTHIA. One of these studies — involving PDE benchmarks — is described in greater detail in the next section.

4.1 Numerical Quadrature

The recommendation problem addressed by this study is to select an method to evaluate a given integral $I = \int_a^b f(x)dx$ so that relative error $\epsilon_r < \theta$ and N_i is minimized; where θ is an user–specified error requirement and N_i is the number of times $f(x)$ is evaluated in $[a, b]$ to yield the desired accuracy. For this testcase of MyPYTHIA, we have utilized 124 routines from the GAMS cross-index [4], where the category H2a is for the evaluation of one dimensional integrals. A collection of 286 test problems and 124 methods were identified. The libraries from which the routines are obtained are QUADPACK, NAG, IMSL, PORT, SLATEC, JCAM and the collected methods of the ACM (TOMS). The test problems were selected so that they exhibit interesting or common features such as smoothness (or its absence), singularities, discontinuities, peaks, and oscillation. Some of the functions were selected so that they satisfy the special considerations on which some methods are designed. For example, routine QDAWO requires that its argument contain a sine or a cosine. Most of the functions are parameterized, which generates families of integrands with similar features and characteristics – this aids in the generalization of the system. For each routine and each applicable integrand, experiments were conducted with varying requirements on

the relative error accuracy ϵ_r. The number of accuracy levels was 10 and the strictest error requirement used was 10^{-8}.

Detailed results are presented in [30]. We briefly summarize that the knowledge discovered through our methodology correspond to accepted general knowledge about numerical integration routines. It was observed, for example, that the adaptive methods use fewer function evaluations to achieve high–accuracy results than their non–adaptive counterparts; conversely, they use more evaluations to meet low–accuracy constraints. A high accuracy adaptive method has been found to be more suitable for an oscillating integrand. This could possibly be due to the fact that in an oscillating function, subdivisions are spread over the entire domain of integration and hence a smaller number of subdivisions are required to achieve a fairly high degree of accuracy. Conversely, integrands with singularities or peaks are more amenable to low and medium accuracy adaptive routines. There are many more such observations, and we have reproduced only the most interesting here. Finally, this implementation has helped identify 'redundant' methods, i.e., methods which perform almost exactly the same for the test functions considered in this work. For example, the rules selecting the methods DPCHIA, DCSITG and DCSPQU contained the same antecedents. DPCHIA evaluates the given integral using piecewise cubic Hermite functions, DCSITG evaluates the integral of a cubic spline and DCSPQU also uses spline interpolation. One interpretation for this result is that quadrature methods using fitted functions often utilize similar mechanisms for obtaining data-dependent break points (for use in Hermite cubics or as knots in spline-based methods).

4.2   PDE Software

The recommendation problem addressed by this case study is to select an method to solve $Lu = f$ on Ω, $Bu = g$ on $\partial\Omega$, such that relative error $\epsilon_r \leq \theta$ and time $t_s \leq T$, where L is a second order, linear elliptic operator, B is a differential operator involving up to first order partial derivatives of u, Ω is a bounded open region in 2–dimensional space, and θ, T are performance criteria constraints. In this study, we restrict ourselves to rectangular domains. Accuracy is measured as the maximum absolute error on the rectangular mesh divided by the maximum absolute value of the PDE solution. Performance studies are conducted and the amount of time required to obtain three levels of accuracy — 10^{-3}, 10^{-4} and 10^{-5} — is tabulated. We use linear least squares approximation to the profile data as this allows us to interpolate between the discrete values of the meshes to specify the size necessary to obtain a specified accuracy. We use the population in [33] of 56 linear second–order elliptic partial differential equations. The primary motivation for developing this population was to aid in the performance evaluation of numerical software. Forty two of the PDEs from this population are parameterized which gives rise to a huge number of PDEs, numbering nearly 200. We use several features of these PDEs to aid in method selection. The principal characteristics used are those of the operator and right side (analytic, constant coefficients, nearly singular behavior, oscillatory, jump discontinuity, singular, peaked, singular derivatives, entire), boundary conditions (as being mixed, homogeneous, Dirichlet, Neumann, constant coefficients, etc.) and those of the domain (unit square, variable square, rectangle, etc.).

This case study is described in greater detail in Section 5. We briefly summarize the results here. The rules discovered confirm the statistically discovered conclusion in [15] that higher order methods are better for elliptic PDEs with singularities. They also confirm the general hypothesis that there is a strong correlation between the order of a method and its efficiency. There were several other interesting inferences drawn. The rule that had the maximum cover from the data was the one which stated that the 9-point FFT method is best for a PDE if the PDE has a Laplacian operator, homogeneous and Dirichlet boundary conditions and discontinuous derivatives on the right side. Finally, an approximate ordering was induced for the overall population. This gave rise to the ordering whose posited ranks correspond most closely to that for Poisson problems which formed the bulk of our population. In overall, the rules from this study performed best method recommendation for 100% of the cases.

4.3 Performance Modeling of Large Scientific Codes

Our final case study involved assessing performance models of Sweep3D, a complex ASCI benchmark for the solution of 3D time-independent, neutron particle transport equations [2]. Available in FORTRAN code in the public domain, it forms the kernel of an ASCI application and involves solving a group of time-independent discrete-ordinate (Sn), neutron particle transport equations on a XYZ cartesian cube [20]. The XYZ geometry is represented by an IJK logically rectangular grid of cells. The angular dependence is handled by discrete angles with a spherical harmonics treatment for the scattering source. The solution involves two steps: the streaming operator is solved by sweeps for each angle and the scattering operator is solved iteratively.

The first recommendation goal is to facilitate the selection of computational parameters of the application/machine pair in order to achieve pre-specified performance goals. For example, in the context of SWEEP3D application, it can be used to obtain the parameters of the underlying algorithm (e.g., grid size, spacing, scattering order, angles, k-blocking factor, convergence test), system (e.g., I/O switches), and the target machine (e.g., processor configuration). Thus, in the case of a fixed application/architecture pair where the parameters are a priori known and there are performance data for various values of the parameters, then MyPYTHIA can be used a) to identify the best pair from the given performance data for different levels of performance, b) to identify the best pair for a specified performance objective (i.e., execution time, cost), and c) to identify the best pair including an estimate of its parameters by specifying certain a priori selected features of the pair.

The second recommendation goal is to automatically classify the various designs/implementations based on their performance data by using well designed benchmarks with specific features and with respect to range values of some performance indicators or features. Here, the user might want to see a classification of the data with respect to performance levels or certain features (i.e., architectural model, shared memory, distributed memory, etc.)

The third goal is to predict the performance of a conceptual design by comparing it with the performance of existing 'similar' designs/implementations and assuming some user defined computational goals and design features.

5 A Complete MyPYTHIA Application

We present a case study that demonstrates how a knowledge engineer can use MyPYTHIA to discover and evaluate the pertinent knowledge needed to solve the software selection problem for a certain population of PDEs.

For this case study we have selected 30 benchmark PDE problems from [15] for which the solutions exhibit singularities or near singularities. Each of the 30 problems is parameterized by 3 sets of equation coefficients and 5 different grid sizes, resulting in a total population of $30 \times 3 \times 5 = 450$ unique PDE problems. The seven methods to be evaluated by this case study of MyPYTHIA for solving the problems are listed in Table 1. Each of the 450 problems represents a subset of the feature space which contains 64 features. These features have been categorized into six classes: features for the PDE operator, the PDE, the right hand side, the solution, the boundary conditions, and the domain. Each class of features is shown in a separate table: Table 2 lists the operator features, Table 3 lists the PDE features, Table 4 lists the right hand side features, Table 5 lists the boundary condition features, Table 6 lists the solution features, and Table 7 lists the domain features.

As described in a previous section, the performance data must be generated by the PDE expert outside of the MyPYTHIA system, and then imported into MyPYTHIA in an XML-like format according to the underlying database schema for the MyPYTHIA performance records. In the rest of this section, we present some of the important principles related to the modeling and analysis of this domain using MyPYTHIA.

The 'feature-measure context' for this case study contains logical references to 30 different analyzer ranking runs. Each analyzer run considers one of the benchmark PDEs across the 3 coefficient parameter sets and 5 grid size parameters. Each run addresses the two measures selected by the PDE expert for analyzing the effectiveness and efficiency of the 7 numerical methods: accuracy (error in the computed solution) and elapsed CPU time in solving

Table 1: Numerical methods to be evaluated by MyPYTHIA.

Program Name	Method Implemented
5-POINT STAR	Ordinary second-order finite differences; Gauss Elimination
P3C1-COLLOCATION	Fourth-order collocation with Hermite bicubics; Gauss Elimination
DYAKANOV CG	Ordinary second-order finite differences; iteration with generalized marching method and conjugate gradient method
DYAKANOV CG-4	Dyakanov CG with Richardson extrapolation to achieve fourth-order accuracy
FFT9(IORDER=2)	Ordinary second-order finite differences and Fast Fourier Transform
FFT9(IORDER=4)	Fourth-order finite difference method and Fast Fourier Transform
FFT9(IORDER=6)	Sixth-order finite difference method and Fast Fourier Transform

Table 2: Symbolic names and descriptions of features for the operator.

Feature Name	Description
opConstCoeff	Constant Coefficient
opGeneral	General Operator
opHelmholtz	$AU_{xx} + CU_{yy} + fU = 0$
opLaplace	$U_{xx} + U_{yy}(+U_{zz}) = f$
opSelfAdjoint	$(PU_x)_x + (QU_y)_y + (RU_z)_z + Sf = g$
opPoisson	$U_{xx} + U_{yy}(+U_{zz}) = 0$
opSmoConst	Constant Operator Smoothness

Table 3: Symbolic names and descriptions of features for the PDE.

Feature Name	Description
pdeAnalytic	PDE problem is analytic
pdeEntire	PDE problem is entire
pdeConstCoeff	PDE problem has constant coefficients
pdeNearlySingular	PDE problem is nearly singular
pdeOscillatory	PDE problem is oscillatory
pdeJump	PDE problem has a jump discontinuity
pdeSingular	PDE problem is singular (infinite)
pdePeaked	PDE problem is parameterized or peaked
pdeCompComplex	PDE problem is computationally complex
pdeSingDeriv	PDE problem has singular derivatives
pdeSmoConst	Smoothness of PDE problem is constant
pdeSmoEntire	Smoothness of PDE problem is entire
pdeSmoAnalytic	Smoothness of PDE problem is analytic
pdeSmoDiscDeriv	Smoothness of PDE problem has discontinuous derivatives
pdeSmoSingular	Smoothness of PDE problem is singular
pdeSmoOscillatory	Smoothness of PDE problem is oscillatory
pdeSmoCompComp	Smoothness of PDE problem is computationally complicated

Table 4: Symbolic names and descriptions of features for the right hand side.

Feature Name	Description
rhsEntire	Right hand side is entire
rhsAnalytic	Right hand side is analytic
rhsSingular	Right hand side is singular (infinite)
rhsSingDeriv	Right hand side has singular derivatives
rhsConstCoeff	Right hand side has constant coefficients
rhsNearlySingular	Right hand side is nearly singular
rhsPeaked	Right hand side is parameterized or peaked
rhsOscillatory	Right hand side is oscillatory
rhsHomogeneous	Right hand side is homogeneous
rhsCompComplex	Right hand side is computationally complex
rhsSmoConst	Smoothness of right hand side is constant
rhsSmoEntire	Smoothness of right hand side is entire
rhsSmoAnalytic	Smoothness of right hand side is analytic
rhsSmoDiscDeriv	Smoothness of right hand side has discontinuous derivatives
rhsSmoSingular	Smoothness of right hand side is singular
rhsSmoOscillatory	Smoothness of right hand side is oscillatory
rhsSmoPeaked	Smoothness of right hand side is peaked
rhsSmoCompComp	Smoothness of right hand side is computationally complicated

Table 5: Symbolic names and descriptions of features for the boundary condition.

Feature Name	Description
bcHomogeneous	$U = 0$ on boundary
bcDirichlet	$AU = f$ on boundary
bcNeumann	$BU_n = f$ on some boundary
bcMixed	$AU + BU_n = f$ on some boundary
bcConstCoeff	Boundary condition has constant coefficients
bcVarCoeff	Boundary condition has non-constant coefficients

Table 6: Symbolic names and descriptions of features for the solution.

Feature Name	Description
solEntire	Solution is Entire
solAnalytic	Solution is Analytic
solSingular	Solution is singular (infinite)
solSingDeriv	Solution has singular derivatives
solOscillatory	Solution is oscillatory
solWaveFront	Solution has a wave front
solBoundLayer	Solution has a boundary layer
solPeaked	Solution is parameterized or peaked
solUnknown	Solution is unknown
solNearSingular	Solution is nearly singular
solVarSmooth	Solution has variable smoothness
solSmoEntire	Smoothness of the solution is entire
solSmoAnalytic	Smoothness of the solution is analytic
solSmoSingDeriv	Smoothness of the solution has singular derivatives
solSmoOscillatory	Smoothness of the solution is oscillatory
solSmoWaveFront	Smoothness of the solution has wave front
solSmoDiscDeriv	Smoothness of the solution has discontinuous derivatives
solSmoSingular	Smoothness of the solution is singular
solSmoBoundLayer	Smoothness of the solution has a boundary layer
solSmoPeak	Smoothness of the solution has a peak
solSmoTabled	Smoothness of the solution is tabled

Table 7: Symbolic names and description of features for the domain.

Feature Name	Description
domVariableSquare	$(a, b) \times (a + x, b + x)$, where x can vary
domUnitSquare	$(0, 1) \times (0, 1)$
domSquare	$(a, b) \times (a + c, b + c)$ where c is a constant
domRectangle	$(a, b) \times (c, d)$
domVariableRectangle	$(x, y) \times (z, u)$

the problem. As each analyzer run is initiated, the values for these two measures are retrieved from the performance database records for the parameterized variations of a specific benchmark problem. Thus, $(3 \times 5) \times 7 = 105$ problem-method instances of elapsed time and error are extracted for each run, and $450 \times 7 = 3150$ problem-method instances are accessed for the complete ranking analysis.

The full list of features to be used in the knowledge discovery process is also identified in the 'feature-measure context.' We have selected 64 features covering the entire suite of classes for the operator, PDE, right hand side, solution, solution, boundary conditions and domain. All of the features are identified by symbolic names and are associated with boolean values, i.e., the 'possible values' assumed by these features listed in the FEATURE recored is true/false. The boolean value indicates whether the named feature belongs to the PDE problem or not.

The 'feature-measure context' controls the entire process of analyzer ranking and rules generation. The analyzer builds performance profiles for each numerical method applied to each unique PDE problem, maintaing the resulting profile data in the database. The profiles are used by the statistical analyzer to rank the methods for each of the 30 benchmarks, in each case using a rank sum analysis applied to the rankings across the parameter variations of a given benchmark problem. After the method ranking for each benchmark record has been determined, the benchmark's features to be considered (according to the 'feature-measure context') are retrieved and a predicate is generated attaching the method ranked first as the `best` method for the corresponding benchmark problem.

The predicates generated in this way, along with the predicates stating the features of each benchmark, are saved. Based on the knowledge discovery process selected by the domain expert, different filters are used by the system to transform the predicate information to an appropriate format for the knowledge discovery tool. MyPYTHIA system utilizes many tools for the knowledge discovery phase, but in this case study we discuss only the `PROGOL` system, because it is the only one that is currently interfaced with MyPYTHIA's web Recommender system.

`PROGOL` [5] induces logical descriptions of relations in a subset of first order predicate logic. This means that the induction process is extremely expensive. For this reason, the concept space to be explored by `PROGOL` must be well suited to the domain under consideration. This becomes obvious if we consider that `PROGOL` immediately fails to induce any rule for the 64-dimensional space of our case study due to resource limits. A solution to this problem, then, is to reduce the high dimensional space of this study. In order to accomplish this, we make use of the *feature subset selection* utility that is offered by the `MLC++` library. The feature subset selection is a wrapper inducer (in the sense that it wraps around a regular inducer) and selects a good subset of features for improved accuracy performance. This topic has been fully covered in [21].

By applying this methodology to the feature space in this case study we end up with only four features that are chosen by the feature subset selection process as the most important for the identification of the classes. The results from the feature subset selection process are illustrated in Figure 6.

The most important information to be acquired from the feature subset selection process is the list of numbers inside the square brackets which specify the index of the features considered as the most important by the program, along with an estimation of the error attained for each one of the subsets of features selected as best each time. The indices of the selected features are: 1, 14, 19, and 21. The indexing scheme assumed by the method considers the first feature as having index 0. For this reason the selected features are: `opGeneral`, `pdeSmoCompComp`, `rhsAnalytic`, and `rhsEntire`. We also observe that the error falls from 70% to 30%. The wrapped inducer in

```
New best node (1 evals) #0[]: error: 70.00% +- 10.48% (0.00% -
100.00%).
   Test Set: 70.00% +- 8.51% [52.12% - 83.34%].  Bias: 0.00% cost: 10
    complexity: 0
...............................................................
New best node (68 evals) #66[1, 14, 21]: error: 36.67% +- 7.78%
     (0.00% - 66.67%).
   Test Set: 33.33% +- 8.75% [19.23% - 51.22%].  Bias: 3.33% cost: 10
    complexity: 3
..............................................................
New best node (132 evals) #86[1, 14, 19, 21]: error: 30.00% +- 5.98%
     (0.00% - 66.67%).
   Test Set: 26.67% +- 8.21% [14.18% - 44.45%].  Bias: 3.33% cost: 10
    complexity: 4
```

Figure 6: Output generated by the feature subset selection process for the feature space of the PDE with (near) singularities case study.

Table 8: Confusion matrix for the wrapped `ID3` inducer.

(a)	(b)	(c)	(d)	(e)	(f)	(g)	Classes
4.00	0.00	0.00	0.00	0.00	0.00	0.00	**(a) 5-point_star_and_band_ge**
1.00	3.00	0.00	0.00	0.00	0.00	0.00	**(b) hermite_collocation_and_band_ge**
0.00	0.00	0.00	0.00	0.00	0.00	0.00	**(c) dyakanov-cg**
0.00	0.00	0.00	6.00	0.00	1.00	2.00	**(d) dyakanov-cg4**
0.00	0.00	0.00	1.00	0.00	0.00	0.00	**(e) fft_9_point_order_2**
0.00	0.00	0.00	0.00	0.00	5.00	1.00	**(f) fft_9_point_order_4**
0.00	0.00	0.00	1.00	0.00	1.00	4.00	**(g) fft_9_point_order_6**

this case is the `ID3` inducer. The final error for the `ID3` inducer using a 10-fold cross validation is only 26.667%. Table 8 shows the confusion matrix for the problem instances used in this case study.

The rule set produced by `PROGOL` along with coverage statistics of each non-trivial (compressed) rule is shown in Figure 7. Each logical rule generated by `PROGOL` is associated with an information compression factor measuring the generalization accuracy of the rule. Its simple formula is $f = p - (c + n + h)$ where p and n are the number of positive and negative examples respectively covered by a specific rule, c is the number of predicates in the body of the clause/rule, and h is the number of predicates which must be added to the clause to complete the relationship between variables in the head of the clause. The information compression factor is used for ordering the rules in the rule base in a decreasing order and values for each one of these rule parameters are shown in Figure 7.

The confusion matrix by using the leave-one-out accuracy estimation technique for the rule set shown in Figure 7 is illustrated in Figure 8. The column labeled A in the contingency table represents the examples that are known to belong to the set of positive examples while the column labeled Ã consists of negative examples. The positive and negative examples, in sequel, are separated into two different categories: the category P, which includes the number of examples that they were predicted as positive and P̃, the examples that were predicted as negative examples. Additional information related to confidence and overall accuracy is also shown in Figure 8.

```
best_method(p2,fft_9_point_order_2).
best_method(p9,fft_9_point_order_6).
best_method(p19,fft_9_point_order_4).
best_method(A,5-point_star_and_band_ge) :- opGeneral_yes(A).
f=278,p=40,n=1,h=0 best_method(A,hermite_collocation_and_band_ge)
:- pdeSmoCompComp_yes(A). f=278,p=40,n=1,h=0
best_method(A,fft_9_point_order_4) :- rhsEntire_yes(A).
f=276,p=50,n=2,h=0 best_method(A,dyakanov-cg4) :- opGeneral_no(A),
pdeSmoCompComp_no(A). f=240,p=80,n=13,h=0
best_method(A,fft_9_point_order_6) :- opGeneral_no(A),
   pdeSmoCompComp_no(A), rhsEntire_no(A).
f=216,p=50,n=10,h=0 [Total number of clauses = 8]
```

Figure 7: Rule set generated by PROGOL rule generation program along with coverage statistics.

```
[Partial accuracy= 149/210] Contingency table=
                    _____A_____~A
              P|        24|        55|        79
               |(    11.3)|(    67.7)|
             ~P|         6|       125|       131
               |(    18.7)|(   112.3)|
                ~~~~~~~~~~~~~~~~~~~~
                         30       180       210
[Overall accuracy= 70.95% +/- 3.13%]
[Chi-square = 24.72] [Without Yates correc-
tion = 26.79]
[Chi-square probability = 0.0000]
```

Figure 8: Contingency matrix representing accuracy information and statistics for the rule set generated for the PDE case study.

6 Discussion

As a recommendation portal, MyPythia enables the complete automation, generation, and maintenance of domain specific recommender systems, but does not neglect human intervention throughout the process. Our modularized system enables a data-driven approach to grid computing, by clearly defining the interfaces at which human support or iteration can take place. Our open database architecture can easily scale with the quantity and variety of information that will be generated by recommendation studies. By facilitating efficient storage and retrieval for data-intensive computing in scientific simulations, MyPYTHIA serves as a complementary approach to other projects in Grid computing that provide decision-support in the form of data management facilities [8, 19, 23, 24, 28]. MyPYTHIA's system design and architecture are novel in that they formalize the scientific software recommendation problem and point in the direction of connecting such services with grid information services and other knowledge-based portals.

References

[1] C.A. Addison, W.H. Enright, P.W. Gaffney, I. Gladwell, and P.M. Hanson. Algorithm 687: A Decision Tree for the Numerical Solution of Initial Value Ordinary Differential Equations. *ACM Transactions on Mathematical Software*, Vol. 17(1):pages 1–10, 1991.

[2] V.S. Adve, R. Bagrodia, J.S. Browne, E. Deelman, A. Dube, E.N. Houstis, J.R. Rice, R. Sakellariou, D.J. Sundaram-Stukel, P.J. Teller, and M.K. Vernon. POEMS: End-to-End Performance Design of Large Parallel Adaptive Computational Systems. *IEEE Transactions on Software Engineering*, Vol. 26(11):pages 1027–1048, November 2000.

[3] D.W. Aha, D. Kibler, and M.K. Albert. Instance-Based Learning Algorithms. *Machine Learning*, Vol. 6(1):pages 37–66, 1991.

[4] R.F. Boisvert. The NIST Guide to Available Mathematical Software. URL: `http://gams.nist.gov`, 1996.

[5] I. Bratko and S. Muggleton. Applications of Inductive Logic Programming. *Communications of the ACM*, Vol. 38(11):pages 65–70, November 1995.

[6] H. Casanova and J. Dongarra. NetSolve: A Network Server for Solving Computational Science Problems. *International Journal of Supercomputer Applications and High Performance Computing*, Vol. 11(3):pages 212–223, Fall 1997.

[7] K.M. Chandy. The Scientist's Infosphere. *IEEE Computational Science & Engineering*, Vol. 3(2):pp. 43–44, Summer 1996.

[8] I. Foster, J. Insley, G. Von Laszewski, C. Kesselman, and M. Thiebaux. Distance Visualization: Data Exploration on the Grid. *IEEE Computer*, Vol. 32(12):pages 36–43, December 1999.

[9] I. Foster and C. Kesselman. Globus: A Metacomputing Infrastructure Toolkit. *International Journal of Supercomputer Applications and High Performance Computing*, Vol. 11(2):pages 115–128, Summer 1997.

[10] I. Foster and C. Kesselman. *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann Publishers, July 1998.

[11] A.S. Grimshaw, A. Ferrari, G. Lindahl, and K. Holcomb. MetaSystems. *Communications of the ACM*, Vol. 41(11):pages 46–55, November 1998.

[12] A.S. Grimshaw, J.B. Weismann, E.A. West, and Ed. C. Loyot Jr. MetaSystems: An Approach Combining Parallel Processing and Heterogeneous Distributed Computing Systems. *Journal of Parallel and Distributed Computing*, Vol. 21(3), June 1994.

[13] M. Hollander and D.A. Wolfe. *Nonparametric Statistical Methods*. John Wiley and Sons, 1973.

[14] E.N. Houstis, A.C. Catlin, J.R. Rice, V.S. Verykios, N. Ramakrishnan, and C.E. Houstis. PYTHIA-II: A Knowledge/Database System for Managing Performance Data and Recommending Scientific Software. *ACM Transactions on Mathematical Software*, Vol. 26(2):pages 227–253, June 2000.

[15] E.N. Houstis and J.R. Rice. High Order Methods for Elliptic Partial Differential Equations with Singularities. *International Journal for Numerical Methods in Engineering*, Vol. 18:pages 737–754, 1982.

[16] E.N. Houstis, V.S. Verykios, A.C Catlin, N. Ramakrishnan, and J.R. Rice. *A Data Mining Environment for Modeling the Performance of Scientific Software*, chapter 21, pages 261–271. Enabling Technologies for Computational Science: Frameworks, Middleware and Environments. Kluwer Academic Publishers, 2000.

[17] M.I. Jordan and C.M. Bishop. Neural networks. In A. Tucker, editor, *The Computer Science and Engineering Handbook*, pages 536–556. CRC Press, 1997.

[18] A. Joshi, S. Weerawarana, N. Ramakrishnan, E.N. Houstis, and J.R. Rice. Neuro–Fuzzy Support for PSEs: A Step Toward the Automated Solution of PDEs. *Special Joint Issue of IEEE Computer & IEEE Computational Science and Engineering*, Vol. 3(1):pages 44–56, 1996.

[19] S. Karin and S. Graham. The High Performance Computing Continuum. *Communications of the ACM*, Vol. 41(11):pages 32–35, November 1998.

[20] K.R. Koch, R.S. Baker, and R.E. Alcouffe. Solution of the First-Order Form of the 3-D Discrete Ordinates Equation on a Massively Parallel Processor. *Transactions of the American Nuclear Society*, Vol. 65(198), 1992.

[21] R. Kohavi. *Wrappers For Performance Enhancement and Oblivious Decision Graphs*. PhD thesis, Department of Computer Science, Stanford University, 1995.

[22] B. LaRose. The Development and Implementation of a Performance Database Server. Technical Report CS-93-195, University of Tennessee, August 1993.

[23] R.W. Moore, C. Baru, R. Marciano, A. Rajasekar, and M. Wan. Data-Intensive Computing. In C. Kesselman and I. Foster, editors, *The Grid: Blueprint for a New Computing Infrastructure*, pages 105–129. Morgan Kaufmann, 1998.

[24] R.W. Moore, T.A. Prince, and M. Ellisman. Data-Intensive Computing and Digital Libraries. *Communications of the ACM*, Vol. 41(11):pages 56–62, November 1998.

[25] J.R. Quinlan. Induction of Decision Trees. *Machine Learning*, Vol. 1(1):pages 81–106, 1986.

[26] J.R. Quinlan. *C4.5: Programs for Machine Learning*. Morgan Kaufmann, San Mateo, CA, 1993.

[27] N. Ramakrishnan. *Recommender Systems for Problem Solving Environments*. PhD thesis, Dept. of Computer Sciences, Purdue University, 1997.

[28] N. Ramakrishnan and A.Y. Grama. Mining Scientific Data. *Advances in Computers*, Vol. 55:pages 119–169, 2001.

[29] N. Ramakrishnan and C.J. Ribbens. Mining and Visualizing Recommendation Spaces for Elliptic PDEs with Continuous Attributes. *ACM Transactions on Mathematical Software*, Vol. 26(2):pages 254–273, June 2000.

[30] N. Ramakrishnan, J.R. Rice, and E.N. Houstis. GAUSS: An Online Algorithm Selection System for Numerical Quadrature. *Advances in Engineering Software*, 2001. to appear.

[31] J.R. Rice. A Metalgorithm for Adaptive Quadrature. *Journal of the ACM*, Vol. 22(1):pages 61–82, 1973.

[32] J.R. Rice. The Algorithm Selection Problem. *Advances in Computers*, Vol. 15:pages 65–118, 1976.

[33] J.R. Rice, E.N. Houstis, and W.R. Dyksen. A Population of Linear, Second Order, Elliptic Partial Differential Equations on Rectangular Domains, Part I. *Mathematics of Computation*, Vol. 36:pp. 475–484, 1981.

[34] V.S. Verykios. *Knowledge Discovery in Scientific Databases*. PhD thesis, Department of Computer Sciences, Purdue University, 1999.

[35] V.S. Verykios, E.N. Houstis, and J.R. Rice. Mining the Performance of Complex Systems. *Proceedings of IEEE International Conference on Information Intelligence and Systems*, pages 606–612, November 1999.

[36] V.S. Verykios, E.N. Houstis, and J.R Rice. A Knowledge Discovery Methodology for the Performance Evaluation of Scientific Software. *Journal of Neural, Parallel & Scientific Computations*, 8(2):115–132, 2000.

[37] V.S. Verykios, E.N. Houstis, L.H. Tsoukalas, and K.N. Pantazopoulos. Automating the Analysis of Option Pricing Algorithms through Intelligent Knowledge Acquisition Approaches. *IEEE Transactions on Systems, Man and Cybernetics: Part A.*, 2001. to appear.

[38] S. Weerawarana, E.N. Houstis, J.R. Rice, A. Joshi, and C.E. Houstis. PYTHIA: A Knowledge Based System to Select Scientific Algorithms. *ACM Transactions on Mathematical Software*, Vol. 22:pages 447–468, 1996.