

E-MaLeS 1.1

Daniel Kühlwein¹, Stephan Schulz², and Josef Urban¹

¹ Radboud University Nijmegen, Nijmegen, Netherlands*

² Technische Universität München, München, Germany

Abstract. Picking the right search strategy is important for the success of automatic theorem provers. E-MaLeS is a meta-system that uses machine learning and strategy scheduling to optimize the performance of the first-order theorem prover E. E-MaLeS applies a kernel-based learning method to predict the run-time of a strategy on a given problem and dynamically constructs a schedule of multiple promising strategies that are tried in sequence on the problem. This approach has significantly improved the performance of E 1.6, resulting in the second place of E-MaLeS 1.1 in the FOF divisions of CASC-J6 and CASC@Turing.

1 Introduction

Automatic theorem provers (ATPs) for first-order logic search for proofs of a conjecture in a potentially infinite space of derivations. Experience has shown that no single search strategy can be expected to perform well over very diverse proof problems. Thus, most theorem provers provide dozens or even hundreds of parameters. For systems based on modern equational calculi, parameters include clause selection schemes, term orderings, inference and reduction rules used, etc.

The theorem prover E [8] uses a language for describing useful combinations of parameters as *strategies*. Such strategies can be automatically evaluated over large problem sets and compiled directly into C source code. Over 200 strategies have been named and evaluated for E 1.6, and the best of them are included in the E source code.

This large number of strategies with associated performance data suggests the use of data-driven methods to estimate how to solve new problems. E is one of the first ATPs that have applied machine learning to strategy selection. Below, we first describe how the *automatic mode* of E is generated using simple analogy-based learning. We then introduce an alternative method based on kernel-based learning that automatically determines a schedule of E strategies. The resulting strategy-scheduling meta-system, E-MaLeS 1.1, outperforms the E's original single-strategy automated mode.

2 Learning of Strategy Selection

Characterizing ATP problems is a hard problem. An ambitious formulation could be for example:

* The authors were supported by the NWO projects “MathWiki a Web-based Collaborative Authoring Environment for Formal Proofs” and “Learning2Reason”.

Given a set of ATP strategies and a large class of ATP problems problems, find a (typically finite) set of efficiently computable problem features that can be used to efficiently partition the set of problems into subclasses that share a common best strategy.

The set of features is in practice suggested by the intuition of the ATP implementer or other domain expert (for example, a mathematician who might know what features could be important for distinguishing different classes of problems in his domain). Some of the features can correspond to precise knowledge about logical calculi, for example, if the problem is Horn, or effectively propositional. Others can express hunches, for example, if there are many or few clauses, symbols, terms, etc. If the set of problems uses symbols consistently (this is the case for recent large-theory corpora created from the Mizar and Isabelle libraries, and for the SUMO and Cyc common-sense ontologies), then the (combinations of) symbols (and derived structures, like terms) are often also relevant for proof search.

This paper presents two learning methods that have been developed to estimate the best strategy for a new problem. Both are based on the results of strategies run on the TPTP problems [12]. The traditional one by Stephan Schulz is an instance of the *case-based reasoning* method. The newer alternative uses *kernel-based* learning. Both methods rely on the TPTP being a sufficiently representative set of ATP problems of different kinds, however the methods can be applied for any sufficiently big corpus of problems (e.g. the MPTP [14]). Both methods require a relevant *feature characterization* of problems and a record of the performance of different strategies for training. Once trained, both provide suggestions for strategies to use on new problems, either a single strategy for the original method, or a schedule of strategies for the newer method.

2.1 E's Feature Characterization

The set of features used in E for problem characterization is listed in Table 1. All features apply to the clausal form of a problem. A clause is called *negative* if it only has negative literals. It is called *positive* if it only has positive literals. A ground clause is a clause that contains no variables. In this setting, we refer to all negative clauses as “goals”, and to all other clauses as “axioms”. Clauses can be *unit* (having only a single literal), *Horn* (having at most one positive literal), or *general* (no constraints on the form). All unit clauses are Horn, and all Horn clauses are general. Goals have no positive literals and are hence always at least Horn.

2.2 E's Automatic Mode

E supports an automatic mode that analyzes the problem and determines all major search parameters (literal selection function, clause selection heuristics, term ordering, and a number of mostly discrete options controlling optional simplifications and preprocessing steps). It has been conservatively extended from the very first implementation in E 0.3 *Castleton*, released in 1999.

Table 1. Problem features used by strategy selection in E

Feature	Description
<code>axioms</code>	Most specific class (unit, Horn, general) describing all axioms
<code>goals</code>	Most specific class (unit, Horn) describing all goals
<code>equality</code>	Problem has no equational literals, some equational literals, or only equational literals
<code>non_ground_units</code>	Number of unit axioms that are not ground
<code>ground_goals</code>	Are all goals ground?
<code>clauses</code>	Number of clauses
<code>literals</code>	Number of literals
<code>term_cells</code>	Number of terms (including subterms)
<code>ground_positive_axioms</code>	Number of positive axioms that are ground
<code>max_fun_arity</code>	Maximal arity of a function or predicate symbol
<code>avg_fun_arity</code>	Average arity of symbols in the problem
<code>sum_fun_arity</code>	Sum of arities of symbols in the problem
<code>clause_max_depth</code>	Maximal clause depth

The automatic mode is based on a static partitioning of the set of all CNF problems into disjoint classes. It is generated in two steps. First, the set of all training examples (typically the set of all current TPTP problems) is classified into disjoint classes using the features listed in Table 1. For the numeric features, threshold values have originally been selected to split the TPTP into 3 or 4 approximately equal subsets on each feature. Over time, these have been manually adapted using trial and error.

Once the classification is fixed, a Python program reads the different classes and a set of test protocols describing the performance of different strategies on all problems from the test set. It assigns to each class one of the strategies that solves the most examples in this class. For *large* classes (arbitrarily defined as having more than 200 problems), it picks the strategy that also is fastest on that class. For small classes, it picks the globally best strategy among those that solve the maximum number of problems. A class with zero solutions by all strategies is assigned the overall best strategy.

2.3 Strategy Scheduling

Strategy learning is currently being used by E’s automatic mode to predict the best strategy for a problem. The predicted strategy is then run for the full time.

An alternative approach (known mainly from Gandalf [13], E-SETHEO [11], and Vampire [7]) is strategy scheduling. The idea is to run different strategies for fractions of the overall time. This can help when the strategies are sufficiently orthogonal (solve different problems), and the problem classification is not good enough to clearly point to one best strategy.

In the simplest setting, the suitable set of strategies is considered independent of the problem features. In that case, given a database of results of many strategies on a large set of problems (TPTP, MPTP, etc.), the goal is to cover the set of solvable problems by as few strategies as possible. This is an NP-complete problem, which however seems to be quite easy (for our instances) for systems like MiniSat++ [10]. For example, for an experiment with covering randomly chosen 400 problems from the MPTP2078 benchmark by 280 E strategies, MiniSat++ can find the minimal cover (9 strategies) in 0.09 s.

We present a learning based method that depending on the problem features predicts the time each strategy needs to solve the problem. The predictions are used to schedule in which order and for how long the strategies should be run.

2.4 E-MaLeS 1.1

E-MaLeS (E Machine Learning of Strategies) uses the kernel-based MOR algorithm, see [1] for details. E-MaLeS is freely available at <http://cs.ru.nl/~kuehlwein/>. E-MaLeS learns a function that predicts the performance for each strategy on each problem. Given the features of a problem defined in Table 1, E-MaLeS predicts for each strategy s how long E running s will need to solve the problem. A similar approach has successfully been used in the SAT community [16].

The MOR algorithm is an instance of kernel-based learning. Kernel-based learning is a machine learning approach that finds (typically non-linear in the features) approximations of the training data by minimizing a *loss function* describing the difference between learned approximation and training data. By mapping the data in a higher-dimensional vector space, kernel methods combine the expressiveness of a high-dimensional function space with the simplicity of linear regression. Intuitively, kernels can be seen as a similarity measure between different data points (in our case problems). Kernel-based methods are among the most successful algorithms applied to various problems from bioinformatics to information retrieval to computer vision [9].

E-MaLeS uses a Gaussian kernel. The similarity measure induced by this kernel can be imagined as a Gaussian distribution around each data point with the width σ of the distribution being an adjustable parameter. To ensure that the learned functions generalize well, a regularization parameter λ is used. Regularization adds an additional term based on the complexity of the learned function to the loss function. The more complex a function, the bigger the penalty. The intuition behind this is that complex function are more likely to overfit. The value of λ determines the weight of the regularization term, with $\lambda = 0$ being equivalent to no regularization.

The values for σ and λ are determined via a 10-fold cross-validation. First, we define logarithmically scaled grids of potential values. The training dataset is then shuffled and divided into two parts using a 70/30 split. For each parameter pair the algorithm trains a function based on the data points in the larger part and evaluates it (i.e. compares the predicted run times with the actual run times) on the data points in the smaller part. This process is repeated 10 times. The

parameter pair with the best average performance (i.e. the minimum average squared difference between the predicted run times and the actual run times) on the smaller set is then used for the final learning. The goal of cross-validation is to estimate the performance of the learned function on unseen data points.

During the learning phase, the input to E-MaLeS is a list of strategies and a list of problems with their features, together with the performances of the strategies on the problems within a fixed time limit (e.g. 300 seconds). The features are first normalized to values between 0 and 1. Ideally, we would have the exact time needed until a proof is found for each problem-strategy pair. Unfortunately, real world limitations restrict us to finite run times – in our case 300 seconds. Problems that were not solved within this time limit are ignored. Note that this leads to different training data for different strategies and a bias towards lower times. Thus, the learned prediction functions are likely to predict a time that is lower than the actual needed time. An alternative to simply ignoring unsolved problems would be to use a large fixed time for each (e.g. 600 seconds). However, in initial experiments this did not show any improvement.

When trying to solve a new problem, E-MaLeS employs a combination of E’s automatic mode and strategy scheduling. First, the automatic mode is run for 60 seconds.¹ If the automatic mode fails to find a proof, the features of the problem are computed and normalized.² For each strategy the time needed to solve the problem is predicted using the prediction function learned during the setup. Since E’s timeout parameter expects seconds, the predicted times are rounded up to next full second. The strategies are then run for their predicted time, starting with the strategy with the smallest predicted time.³ If the sum of the rounded predicted times is less than the total time given, the remaining free time is spread equally over all strategies.

3 Results

Both E-MaLeS 1.1 and E 1.6 competed at CASC@Turing and CASC-J6. In both competition, E-MaLeS 1.1 won the second place in the FOF division, solving more problems than E 1.6. The results are shown in Tables 2 and 3

In the FOF division of CASC@Turing, E-MaLeS solved 4.6% more problems than E. If we only compare the results of the new problems, E-MaLeS solved 14.5% more problems. The results for FOF division of CASC-J6 are similar, with E-MaLeS solving 4% more problems than E. On the new problems, E-MaLeS solved 13.2% more problems than E.

¹ Running the auto mode first allows us to reduce the number of training examples of the machine learning algorithm. We only learn from problems that cannot be solved by the automatic mode within 60 seconds. The reason for this is that the learning time of the algorithm is cubic in the number of training examples which makes learning from all examples (the whole TPTP library) infeasible.

² Since the normalization function is defined during setup, the normalized features of new problems may fall out of the $[0, 1]$ interval.

³ A possible improvement would be to take not only the predicted times, but also the orthogonality (difference in problems solved between strategies) of strategies into account. First experiments with this approach were promising.

Table 2. Results for the CASC@Turing problems

System	All Problems	New Problems
E 1.6	378/500	73/97
E-MaLeS 1.1	401/500	87/97

Table 3. Results for the CASC-J6 problems

System	All Problems	New Problems
E 1.6	359/450	50/68
E-MaLeS 1.1	377/450	59/68

A possible explanation of the discrepancy between old and new problems solved is that E 1.6 is overspecialized for the old problems. Cross-validation and regularization helps E-MaLeS to combat such overfitting.

E-MaLeS also competed in the LTB (Large Theory Batch) division of CASC-J6 and placed fourth with 83 problems solved after E 1.6 with 87 solved problems. The LTB division contains problems from large theories, namely problems exported from Isabelle, Mizar and SUMO. Unlike in the FOF division, the ATPs may use all cores of the machine. The LTB version of E-MaLeS uses the extra time and all available cores to run more strategies but is in no other way optimized. E 1.6, on the other hand, makes use of the batch structure to avoid repeated parsing of the large background theories, and also makes better use of the various SInE strategies that heuristically select relevant premises from the large theories.

4 Future Work

There are several ways to improve the current algorithm. Extracting features based on the FOF instead of the CNF representation of a formula could lead to better learning performance. Using different learning algorithms might allow us to run E-MaLeS without relying on E’s automatic mode.

Strategies themselves are just combinations of many ATP parameters. An interesting application of machine learning to ATP is to develop new strategies by searching for such good parameter combinations systematically (by methods like hill-climbing) on a large corpus of problems. A related task is to produce a set of strategies that are highly orthogonal, i.e., that solve very different problems, so that their collective coverage is high.

We could also learn strategies based on conjecture features (like symbols) in consistently-named corpora like MPTP, instead of just using abstract features like on the TPTP problems.

The methods that we develop can probably be used for any ATP, and it would be interesting to see how such learning and optimization works for systems like Vampire [7], Z3 [6], and iProver [5]. Meta-systems like Isabelle/Sledgehammer [4,3] and MaLAREa [15] could then use this deeper knowledge about ATPs to attack new conjectures with the strongest possible combinations of systems and strategies.

References

1. Alama, J., Heskes, T., Kühlwein, D., Tsvitsivadze, E., Urban, J.: Premise Selection for Mathematics by Corpus Analysis and Kernel Methods. CoRR abs/1108.3446 (2011); Accepted to JAR
2. Armando, A., Baumgartner, P., Dowek, G. (eds.): IJCAR 2008. LNCS (LNAI), vol. 5195. Springer, Heidelberg (2008)
3. Blanchette, J.C., Böhme, S., Paulson, L.C.: Extending sledgehammer with SMT solvers. In: Bjørner, N., Sofronie-Stokkermans, V. (eds.) CADE 2011. LNCS, vol. 6803, pp. 116–130. Springer, Heidelberg (2011)
4. Blanchette, J.C., Bulwahn, L., Nipkow, T.: Automatic Proof and Disproof in Isabelle/HOL. In: Tinelli, C., Sofronie-Stokkermans, V. (eds.) FroCoS 2011. LNCS, vol. 6989, pp. 12–27. Springer, Heidelberg (2011)
5. Korovin, K.: iProver - An Instantiation-Based Theorem Prover for First-Order Logic (System Description). Armando et al. [2], 292–298
6. de Moura, L., Bjørner, N.S.: Z3: An Efficient SMT Solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 337–340. Springer, Heidelberg (2008)
7. Riazanov, A., Voronkov, A.: The design and implementation of VAMPIRE. AI Commun. 15(2-3), 91–110 (2002)
8. Schulz, S.: E - A Brainiac Theorem Prover. AI Commun. 15(2-3), 111–126 (2002)
9. Shawe-Taylor, J., Cristianini, N.: Kernel Methods for Pattern Analysis. Cambridge University Press, New York (2004)
10. Sörensson, N., Eén, N.: MiniSat 2.1 and MiniSat++ 1.0 SAT Race, Editions Technical Report, Chalmers University of Technology, Sweden (2008)
11. Stenz, G., Wolf, A.: E-SETHEO: An Automated³ Theorem Prover – System Abstract. In: Dyckhoff, R. (ed.) TABLEAUX 2000. LNCS, vol. 1847, pp. 436–440. Springer, Heidelberg (2000)
12. Sutcliffe, G.: The TPTP Problem Library and Associated Infrastructure: The FOF and CNF Parts, v3.5.0. Journal of Automated Reasoning 43(4), 337–362 (2009)
13. Tammet, T.: Gandalf. Journal of Automated Reasoning 18, 199–204 (1997)
14. Urban, J.: MPTP 0.2: Design, Implementation, and Initial Experiments. J. Autom. Reasoning 37(1-2), 21–43 (2006)
15. Urban, J., Sutcliffe, G., Pudlák, P., Vyskocil, J.: MaLARea SG1- Machine Learner for Automated Reasoning with Semantic Guidance. In: Armando et al. [2], pp. 441–456
16. Xu, L., Hutter, F., Hoos, H.H., Leyton-Brown, K.: SATzilla: portfolio-based algorithm selection for SAT. J. Artif. Int. Res. 32(1), 565–606 (2008)