

Algorithmically generating new algebraic features of polynomial systems for machine learning

Dorian Florescu
Dorian.Florescu@coventry.ac.uk

Matthew England
Matthew.England@coventry.ac.uk

Faculty of Engineering, Environment and Computing,
Coventry University, Coventry, CV1 5FB, UK

Abstract

There are a variety of choices to be made in both computer algebra systems (CASs) and satisfiability modulo theory (SMT) solvers which can impact performance without affecting mathematical correctness. Such choices are candidates for machine learning (ML) approaches, however, there are difficulties in applying standard ML techniques, such as the efficient identification of ML features from input data which is typically a polynomial system.

Our focus is selecting the variable ordering for cylindrical algebraic decomposition (CAD), an important algorithm implemented in several CASs, and now also SMT-solvers. We created a framework to describe all the previously identified ML features for the problem and then enumerated all options in this framework to automatically generate many more features. We validate the usefulness of these with an experiment which shows that an ML choice for CAD variable ordering is superior to those made by human created heuristics, and further improved with these additional features. This technique of feature generation could be useful for other choices related to CAD, or even choices for other algorithms in CASs / SMT-solvers with polynomial systems as input.

1 Introduction

Machine Learning (ML), that is statistical techniques to give computer systems the ability to *learn* rules from data, is a topic that has found great success in a diverse range of fields over recent years. ML is most attractive when the underlying functional relationship to be modelled is complex or not well understood. Hence ML has yet to make a large impact in the fields which form SC^2 , Symbolic Computation and Satisfiability Checking [1], since these prize mathematical correctness and seek to understand underlying functional relationships. However, as most developers would acknowledge, our software usually comes with a range of choices which, while having no effect on the correctness of the end result, could have a great effect on the resources required to find it. These choices range from the low level (in what order to perform a search that may terminate early) to the high (which of a set of competing exact algorithms to use for this problem instance). In making such choices we may be faced with decisions where relationships are not fully understood, but are not the key object of study.

Copyright © by the paper's authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).
In: A. Editor, B. Coeditor (eds.): Proceedings of the XYZ Workshop, Location, Country, DD-MMM-YYYY, published at <http://ceur-ws.org>

In practice such choices may be made by man-made heuristics based on some experimentation (e.g. [18]) or *magic constants* where crossing a single threshold changes system behaviour [11]. It is likely that many of these decisions could be improved by allowing learning algorithms to analyse the data. The broad topic of this paper is ML for algorithm choices where the input is a set of polynomials, which encompasses a variety of tools in computer algebra systems and the SMT theory of [Quantifier Free] Non-Linear Real Arithmetic, [QF]NRA.

There has been little research on the use of ML in computer algebra: only [33] on identifying efficient multi-variate Horner schemes; [28] [27] [24] on the topic of CAD variable ordering choice; [26], [27] on the question of whether to precondition CAD with Groebner Bases; and [31] on deciding the order of sub-formulae solving for a QE procedure. Within SMT there has been significant work on the Boolean logic side e.g. the portfolio SAT solver SATZILLA [46] and MAPLESAT [34] which views solver branching as an optimisation problem. However there is little work on the use of ML to choose or optimise theory solvers.

We note that other fields of mathematical software are ahead in the use of ML, most notably the automated reasoning community (see e.g. [43], [32], [7], or the brief survey in [19]).

1.1 Difficulties with ML for problems in NRA

There are difficulties in applying standard ML techniques to problems in NRA. One is the lack of sufficiently large datasets, which is addressed only partially by the SMT-LIB. The experiment in [26] found the [QF]NRA sections of the SMT-LIB too uniform, and had to resort to random generated examples (although the state of benchmarking in computer algebra is far worse [22]). There have been improvements since then, with the benchmarks increasing both in number and diversity of underlying application. For example, there are now problems arising from biology [4], [23] and economics [38], [39].

Another difficulty is the identification of suitable features from the input with which to train the ML models. There are some obvious candidates concerning the size and degrees of polynomials, and the distribution of variables. However, this provides a starting set (i.e. before any feature selection takes place) that is small in comparison to other machine learning applications. The main focus of this paper is to introduce a method to automatically (and cheaply) generate further features for ML from polynomial systems.

1.2 Contribution and plan

Our main contributions are the new feature generation approach described in Section 3 and the validation of its use in the experiments described in Sections 4–5. The experiments are for the choice of variable ordering for cylindrical algebraic decomposition, a topic whose background we first present in Section 2, but we emphasise that the techniques may be applicable more broadly.

2 Background on variable ordering for CAD

2.1 Cylindrical algebraic decomposition

A *Cylindrical Algebraic Decomposition* (CAD) is a *decomposition* of ordered \mathbb{R}^n space into cells arranged *cylindrically*: the projections of any pair of cells with respect to the variable ordering are either equal or disjoint. The projections form an induced CAD of the lower dimensional space. The cells are (semi)-algebraic meaning each can be described with a finite sequence of polynomial constraints.

A CAD is produced to be *truth-invariant* for a logical formula (so the formula is either true or false on each cell). Such a decomposition can then be used to perform Quantifier Elimination (QE) over the reals, i.e. given a quantified Tarski formula find an equivalent quantifier free formula over the reals. For example, QE would transform $\exists x, ax^2 + bx + c = 0 \wedge a \neq 0$ to the equivalent unquantified statement $b^2 - 4ac \geq 0$. A CAD over the (x, a, b, c) -space could be used to ascertain this, so long as the variable ordering ensured that there was an induced CAD of (a, b, c) -space. We test one sample point per cell and construct a quantifier free formula from the relevant semi-algebraic cell descriptions.

CAD was introduced by Collins in 1975 [15] and works relative to a set of polynomials. Collins' CAD produces a decomposition so that each polynomial has constant sign on each cell (thus truth-invariant for any formula built with those polynomials). The algorithm first projects the polynomials into smaller and smaller dimensions; and then uses these to lift – to incrementally build decompositions of larger and larger spaces according to the polynomials at that level. For further details on CAD see for example the collection [12].

QE has numerous applications throughout science and engineering [42]. Our work also speeds up independent applications of CAD, such as reasoning with multi-valued functions [17] or motion planning [45].

2.2 Variable ordering

The definition of cylindricity and both stages of the algorithm are relative to an ordering of the variables. For example, given polynomials in variables ordered as $x_n \succ x_{n-1} \succ \dots, \succ x_2 \succ x_1$ we first project away x_n and so on until we are left with polynomials univariate in x_1 . We then start lifting by decomposing the x_1 -axis, and then the (x_1, x_2) -plane and so on. The cylindricity condition refers to projections of cells in \mathbb{R}^n onto a space (x_1, \dots, x_m) where $m < n$. There have been numerous advances to CAD since its inception: new projection schemes [35], [37]; partial construction [16], [44]; symbolic-numeric lifting [41], [29]; adapting to the Boolean structure [5], [20]; and adaptations for SMT [30], [9]. However, the need for a fixed variable ordering remains.

Depending on the application, the variable ordering may be determined, constrained, or free. QE, requires that quantified variables are eliminated first and that variables are eliminated in the order in which they are quantified so there is only partial freedom. The discriminant in the example from Section 2.1 could have been found with any CAD which eliminates x first. A CAD for the quadratic polynomial under ordering $a \prec b \prec c$ has only 27 cells, but 115 are required for the reverse ordering.

Since we can switch the order of quantified variables in a statement when the quantifier is the same, we also have some choice on the ordering of quantified variables. For example, a QE problem of the form $\exists x \exists y \forall a \phi(x, y, a)$ could be solved by a CAD under either ordering $x \succ y \succ a$ or ordering $y \succ x \succ a$. In the SMT context there is only a single existentially quantified block and so there is a completely free choice of ordering.

The choice of variable ordering can have a great effect on the time / memory use of CAD, and the number of cells in the output. Brown and Davenport presented a class of problems in which one variable ordering gave output of double exponential complexity in the number of variables and another output of a constant size [10].

2.3 Prior work on choosing the variable ordering

Heuristics have been developed to choose a variable ordering, with Dolzmann *et al.* [18] giving the best known study. After analysing a variety of metrics they proposed a heuristic, `sotd`, which constructs the full set of projection polynomials for each permitted ordering and selects the ordering whose corresponding set has the lowest sum of total degrees for each of the monomials in each of the polynomials. The second author demonstrated examples for which that heuristic could be misled in [6]; and then later showed that tailoring to an implementation could improve performance [21]. These heuristics all involved potentially costly projection operations on the input polynomials.

In [28] the second author of the present paper collaborated to use a support vector machine to choose which of three human made heuristics to believe when picking the variable ordering, based only on simple features of the input polynomials. The experiments identified substantial subclasses on which each of the three heuristics made the best decision, and demonstrated that the machine learned choice did significantly better than any one heuristic overall. This work was picked up again in [24] by the present authors, where ML was used to predict directly the variable ordering for CAD, leading to the shortest computing time, with experiments conducted for four different ML models.

Both [28] and [24] used a set of 11 human identified features. These did lead to good performance of the models, with ML outperforming the prior human created heuristics, but a starting set of 11 features is relatively small for ML and so we hypothesise that identifying more would improve the results.

3 Generating new features algorithmically

3.1 Existing features for sets of polynomials

An early heuristic for the choice of CAD variable ordering is that of Brown [8], which chooses a variable ordering according to the following criteria, starting with the first and breaking ties with successive ones.

- (1) Eliminate a variable first if it appears with the lowest overall degree in the input.
- (2) For each variable calculate the maximum total degree for the set of terms in the input in which it occurs. Eliminate first the variable for which this is lowest.
- (3) Eliminate a variable first if there is a smaller number of terms in the input which contain the variable.

Despite being computationally cheaper than the `sotd` heuristic (because the latter performs projections before measuring degrees) experiments in [28] suggested this simpler measure actually performs slightly better, although

Table 1: Features used by ML in [28] to choose the ordering of 3 variables for CAD.

#	Description	In Framework
1	Number of polynomials	P
2	Maximum total degree of polynomials	$\max_{m,p} (\sum_v d_v^{m,p})$
3	Maximum degree of x_1 among all polynomials	$\max_{m,p} d_1^{m,p}$
4	Maximum degree of x_2 among all polynomials	$\max_{m,p} d_2^{m,p}$
5	Maximum degree of x_3 among all polynomials	$\max_{m,p} d_3^{m,p}$
6	Proportion of x_1 occurring in polynomials	$\text{av}_p (\text{sgn} (\sum_m d_1^{m,p}))$
7	Proportion of x_2 occurring in polynomials	$\text{av}_p (\text{sgn} (\sum_m d_2^{m,p}))$
8	Proportion of x_3 occurring in polynomials	$\text{av}_p (\text{sgn} (\sum_m d_3^{m,p}))$
9	Proportion of x_1 occurring in monomials	$\text{av}_{m,p} (\text{sgn} (d_1^{m,p}))$
10	Proportion of x_2 occurring in monomials	$\text{av}_{m,p} (\text{sgn} (d_2^{m,p}))$
11	Proportion of x_3 occurring in monomials	$\text{av}_{m,p} (\text{sgn} (d_3^{m,p}))$

the key message from those experiments is that there were substantial subsets of problems for which each human-made heuristic made a better choice than the others.

The **Brown** heuristic inspired almost all the features used by the authors of [28], [24] to perform ML for CAD variable ordering, with the full set of 11 features listed in Table 1 (column 3 will be explained later).

3.2 A new framework for generating polynomial features

Our new feature generation procedure is based on the observation that all the measurements taken by the **Brown** heuristic, and all those features used in [28], [24] can be formalised mathematically using a small number of functions. For simplicity, the following discussion will be restricted to polynomials of 3 variables as these were used in the following experiments, but everything generalises in an obvious way to n variables.

Let a problem instance \mathbf{Pr} be defined by a set of P polynomials

$$\mathbf{Pr} = \{\mathcal{P}_p \mid p = 1, \dots, P\}. \quad (1)$$

This is the typical input for producing a sign-invariant CAD. Of course, any problem instance consisting of a logical formula whose atoms are polynomial sign conditions can also have such a set extracted.

In the following we define the notation for polynomial variables and coefficients that will be used throughout the manuscript. Each polynomial with index p , for $p = 1, \dots, P$ contains a different number of monomials, which will be labelled with index m , where $m = 1, \dots, M_p$ and M_p denotes the number of monomials in polynomial p . We note that these are just labels and are not setting an ordering themselves. The degrees corresponding to each of the variables x_1, x_2, x_3 are a function of m and p . These need to be explicitly labelled in order to allow a rigorous definition of our proposed procedure of feature generation.

We can now write each polynomial as

$$\mathcal{P}_p = \sum_{m=1}^{M_p} c^{m,p} \cdot x_1^{d_1^{m,p}} x_2^{d_2^{m,p}} x_3^{d_3^{m,p}}, p = 1, \dots, P. \quad (2)$$

Here, x_v represents the polynomial variables ($v = 1, 2, 3$). Thus for each monomial in each polynomial there is a tuple (m, p) of positive integers that label it. Then in turn we denote by $d_v^{m,p}$ the degree of variable x_v in that monomial, and by $c^{m,p}$ the constant coefficient, i.e., tuple superscripts are giving a label for a monomial in a problem instance. The original indices are simply a labelling and not an ordering of the variables x_1, x_2, x_3 .

Therefore, any one of our problem instances \mathbf{Pr} is uniquely represented by a set of sets

$$\mathbb{S}_{\mathbf{Pr}} = \{ \{ [c^{m,p}, (d_1^{m,p}, d_2^{m,p}, d_3^{m,p})] \mid m = 1, \dots, M_p \} \mid p = 1, \dots, P \}. \quad (3)$$

Observe that each of Brown's measures can now be formalised as a vector of features for choosing a variable.

- (1) Overall degree in the input of a variable: $\max_{m,p} d_v^{m,p}$.

(2) Maximum total degree of those terms in the input in which a variable occurs:

$$\max_{m,p} \text{sgn}(d_v^{m,p}) \cdot (d_1^{m,p} + d_2^{m,p} + d_3^{m,p}).$$

(3) Number of terms in the input which contain the variable: $\sum_{m,p} \text{sgn}(d_v^{m,p})$.

In the latter two we use the sign function to discriminate between monomials which contain a variable (sign of degree is positive) and those which do not (sign of degree is zero). Of course the sign of the degree is never negative.

Define now also the averaging functions

$$\text{av}_m \triangleq \frac{1}{M_p} \sum_m, \quad \text{av}_p \triangleq \frac{1}{P} \sum_p, \quad \text{av}_{m,p} \triangleq \frac{1}{P} \sum_p \frac{1}{M_p} \sum_m.$$

Then all the features in Table 1 can be formalised similarly to Brown's metrics, as shown in the third column of Table 1. We can place all of these scalars into a single framework:

$$f(\mathbf{Pr}) = (g_4 \circ g_3 \circ g_2 \circ g_1 \circ h^{m,p})(\mathbf{Pr}), \quad (4)$$

where

$$h^{m,p}(\mathbf{Pr}) \in \{d_v^{m,p}, \text{sgn}(d_v^{m,p}) \cdot (\sum_{v'} d_{v'}^{m,p}) \mid v = 1, 2, 3\}$$

and g_1, g_2, g_3 , and g_4 are all taken from the set

$$\{\max_p, \max_m, \max_{m,p}, \max_0, \sum_p, \sum_m, \sum_{m,p}, \sum_0, \text{av}_p, \text{av}_m, \text{av}_{m,p}, \text{av}_0, \text{sgn}, \text{sgn}_0\}.$$

In the above set $\max_0, \sum_0, \text{av}_0$ and sgn_0 are all equal to the identity function.

For example, let $\mathbf{Pr} = \{x_1^2 x_2 - x_3, x_1 x_2^4 x_3^2 + x_1 x_3\}$. If $m = 1, p = 2$, then $h^{1,2}(\mathbf{Pr}) \in \{d_v^{1,2}, \text{sgn}(d_v^{1,2}) \cdot (\sum_{v'} d_{v'}^{1,2}) \mid v = 1, 2, 3\} = \{1, 1 \cdot 7, 4, 4 \cdot 7, 2, 2 \cdot 7\}$.

3.3 Generating additional features

We will thus consider deriving all of the other features which fall into this framework, but to do so we must first impose a number of rules.

1. The functions g_1, g_2, g_3, g_4 must all belong to distinct categories of function, i.e. one each of \max , \sum , av , and sgn .
2. Exactly one of the functions g_1, g_2, g_3, g_4 is computed over p and exactly one is computed over m (it may be the same one).
3. The computation over p is always performed by a function g_i with an index i greater or equal to that of the function computing over m .

The expression of $f(\mathbf{Pr})$ can be interpreted as follows. The values $h^{m,p}(\mathbf{Pr})$ are functions of variables m and p . Each of the functions g_1, g_2, g_3, g_4 either leave the function unchanged, or they turn it into a function of fewer variables (first into a function of p , and then into a scalar value, representing the ML feature).

The rules above are justified as follows. Rule 1 reduces the redundancy in the feature set. Rules 2 and 3 guarantee that the feature $f_v(\mathbf{Pr})$ is well defined and is a scalar number. In particular, Rule 3 is necessary because the computation over the terms in a polynomial is dependent on their number, which is not the same for all polynomials.

The final set $\{f^{(1)}(\mathbf{Pr}), \dots, f^{(N_f)}(\mathbf{Pr})\}$ has size $N_f = 1728$ for a problem with 3 variables. This number is attained as follows: we have 12 possible distributions of indexes to the functions g_1, \dots, g_4 as shown in Table 2; then $4!$ possible orderings of those functions; and 6 possible choices for h . So $4! \cdot 6 \cdot 12 = 1728$.

However, many of these features will be identical (e.g. due to a different placement of the identify function). We do not identify these manually now: that task is trivial for a given dataset, but substantial to do in generality.

4 Machine learning experiment with the new features

We now describe a ML experiment to choose the variable ordering for cylindrical algebraic decomposition. The methodology here is similar to that in our recent paper [24] except for the addition of the extra features from Section 3. A more detailed discussion of the methodology can be found in [24].

Table 2: Possible distributions of indices to the function classes in feature framework.

max	av	sum	sgn	max	av	sum	sgn
p, m	0	0	0	p, m	0	0	1
p	m	0	0	p	m	0	1
p	0	m	0	p	0	0	1
0	p, m	0	0	0	p, m	0	1
0	p	m	0	0	p	m	1
0	0	p, m	0	0	0	p, m	1

4.1 Problem set

We use the `nlsat` dataset¹ produced to evaluate the work in [30], thus the problems are all fully existentially quantified. Although there are CAD algorithms that reduce what is being computed based on the quantifiers in the input (most notably via Partial CAD [16]), the conclusions drawn are likely to be applicable outside of the SAT context. We use the 6117 problems with 3 variables from this database, so each has a choice of six different variable orderings. We extracted only the polynomials involved, and randomly divided into two datasets for training (4612) and testing (1505). Only the former was used to tune the parameters of the ML models.

4.2 Software

We used the CAD routine `CylindricalAlgebraicDecompose` which is part of the `RegularChains` Library for MAPLE. This algorithm builds decompositions first of n -dimensional complex space before refining to a CAD of \mathbb{R}^n [14], [13], [3]. We ran the code in Maple 2018 but used an updated version of the `RegularChains` Library (<http://www.regularchains.org>). Training and evaluation of the ML models was done using the `scikit-learn` package [40] v0.20.2 for Python 2.7. The features for ML were extracted using code written in the `sympy` package v1.3 for Python 2.7, as was the `Brown` heuristic. The `sotd` heuristic was implemented in MAPLE as part of the `ProjectionCAD` package [25].

4.3 Timings

CAD construction was timed in a Maple script that was called separately from Python for each CAD (to avoid Maple’s caching of results). The target variable ordering for ML was defined as the one that minimises the computing time for a given problem. All CAD function calls included a time limit. For the training dataset an initial time limit of 4 seconds was used, doubled incrementally if all orderings timed out, until CAD completed for at least one ordering (a target variable ordering could be assigned for all problems using time limits no bigger than 64 seconds). The problems in the testing dataset were processed with a larger time limit of 128 seconds for all orderings (timeout set as 128s).

4.4 Feature simplification

When computed on a set of problems $\{\mathbf{Pr}_1, \dots, \mathbf{Pr}_N\}$, some of the features $f^{(i)}$ turn out to be constant, i.e. $f^{(i)}(\mathbf{Pr}_1) = f^{(i)}(\mathbf{Pr}_2) = \dots = f^{(i)}(\mathbf{Pr}_N)$. Such features will have no benefit for ML and are removed. Further, other features may be repetitive, i.e. $f^{(i)}(\mathbf{Pr}_n) = f^{(j)}(\mathbf{Pr}_n), \forall n = 1, \dots, N$. This repetition may represent a mathematical equality, or just be the case of the given dataset. Either way, they are merged into a single feature for the experiment. After these steps, we are left with 78 features for our dataset: so while a large majority were redundant, we still have seven times those available in [28], [24].

4.5 Feature selection

Feature selection was performed with the training dataset to see if any features were redundant for the ML. We chose the Analysis of Variance (ANOVA) F-value to determine the importance of each feature for the classification task. Other choices we considered were unsuitable for our problem, e.g. the mutual information based selection requires very large amounts of data. Each problem is assigned a target ordering, or class $c = 1, \dots, C$, where $C = 6$. Let $\mathbf{Pr}_{c,n}$ denote problem number n from the training dataset that is assigned class number c , $c = 1, \dots, C$ and $n = 1, \dots, N_c$, where N_c denotes the number of problems that are assigned class c . Thus $\sum_{c=1}^C N_c = N$.

¹Freely available from <http://cs.nyu.edu/~dejan/nonlinear/>

The F-value for feature number i is computed as follows [36].

$$F_i = \frac{\frac{1}{C-1} \sum_{c=1}^C N_c \left(\bar{f}_c^{(i)} - \bar{f}^{(i)} \right)^2}{\frac{1}{N-C} \sum_{c=1}^C \sum_{n=1}^{N_c} \left(f^{(i)}(\mathbf{Pr}_{c,n}) - \bar{f}_c^{(i)} \right)^2}, \quad (5)$$

where $\bar{f}_c^{(i)}$ is the sample mean in class c , and $\bar{f}^{(i)}$ the overall mean of the data:

$$\bar{f}_c^{(i)} = \frac{1}{N_c} \sum_{n=1}^{N_c} f^{(i)}(\mathbf{Pr}_{c,n}), \quad \bar{f}^{(i)} = \frac{1}{N} \sum_{c=1}^C \sum_{n=1}^{N_c} f^{(i)}(\mathbf{Pr}_{c,n}).$$

The numerator in (5) represents the *between-class variability* or *explained variance* and the denominator the *within-class variability* or *unexplained variance*. The three features with the highest F-values were:

$$\begin{aligned} f_{65}(\mathbf{Pr}) &= \max_0 \text{av}_{m,p} \sum_0 \text{sign}(d_2^{m,p}) = \frac{1}{P} \sum_p \frac{1}{M_p} \sum_m \text{sign}(d_2^{m,p}), \\ f_{46}(\mathbf{Pr}) &= \max_0 \sum_p \text{av}_m \text{sign}(d_2^{m,p}) \cdot (\sum_{v'} d_{v'}^{m,p}) \\ &= \sum_p \frac{1}{M_p} \sum_m \text{sign}(d_2^{m,p}) \cdot (\sum_{v'} d_{v'}^{m,p}), \\ f_{76}(\mathbf{Pr}) &= \text{av}_0 \sum_p \max_m \text{sign}(d_2^{m,p}) \cdot (\sum_{v'} d_{v'}^{m,p}) \\ &= \sum_p \max_m \text{sign}(d_2^{m,p}) \cdot (\sum_{v'} d_{v'}^{m,p}). \end{aligned}$$

The new features may be translated back into natural language. For example, feature 65 is the proportion of monomials containing variable x_2 , averaged across all polynomials; feature 46 the sum of the degrees of the variables in all monomials containing variable x_2 , averaged across all monomials and summed across all polynomials; and feature 76 the maximum sum of the degrees of the variables in all monomials containing variable x_2 , summed across all polynomials.

Feature selection did not suggest to remove any features (they all contributed meaningful information), so we proceed with our experiment using all 78.

4.6 ML models

Four of the most commonly used deterministic ML models were tuned on the training data (for details on the methods see e.g. the textbook [2]). We also give an overview of each in [24].

- The K-Nearest Neighbours (KNN) classifier [2, §2.5].
- The Multi-Layer Perceptron (MLP) classifier [2, §2.5].
- The Decision Tree (DT) classifier [2, §14.4].
- The Support Vector Machine (SVM) classifier with RBF kernel [2, §6.3].

Each model was trained using grid search 5-fold cross-validation, i.e. the set was randomly divided into 5 and each possible combination of 4 parts was used to tune the model parameters, leaving the last part for fitting the hyperparameters with cross-validation, by optimising the average F-score. Grid searches were performed for an initially large range for each hyperparameter; then gradually decreased to home into optimal values. This lasted from a few seconds for simpler models like KNN to a few minutes for more complex models like MLP. The optimal hyperparameters selected during cross-validation are in Table 3.

4.7 Comparing with human made heuristics

The ML approaches were compared in terms of prediction accuracy and resulting CAD computing time against the two best known human constructed heuristics **Brown** [8] and **sotd** [18] as discussed earlier. Unlike the ML, these can end up predicting several variable orderings (i.e. when they cannot discriminate). In practice if this were to happen the heuristic would select one randomly (or perhaps lexicographically), however that final pick is not meaningful. To accommodate this, for each problem, the prediction accuracy of such a heuristic is judged to be the percentage of its predicted variable orderings that are also target orderings. The average of this percentage over all problems in the testing dataset represents the prediction accuracy. Similarly, the computing time for such methods was assessed as the average computing time over all predicted orderings, and it is this that is summed up for all problems in the testing dataset.

Table 3: The ML hyperparameters used following optimisation on the training dataset.

Model	Hyperparameter	Value
Decision Tree	Criterion	Gini impurity
	Maximum tree depth	17
K-Nearest Neighbours	Train instances weighting	Inversely proportional to distance
	Algorithm	Ball Tree
Support Vector Machine	Regularization parameter C	316
	Kernel	Radial basis function
	γ	0.08
Multi-Layer Perceptron	Tolerance for stopping criterion	0.0316
	Hidden layer size	18
	Activation function	Hyperbolic tangent
	Algorithm	Quasi-Newton based optimiser
	Regularization parameter α	$5 \cdot 10^{-5}$

5 Experimental Results

The results are presented in Table 4. We compare the four ML models on accuracy, defined as the percentage of problems where they selected the optimum ordering, and also the total computation time (in seconds) for solving all the problems with their chosen orderings. The first two rows of the top table reproduce the results of [24] which used only the 11 features from Table 1, while the latter two rows are the results from the new experiment in the present paper which has 78 features. We also compare with the two human constructed heuristics and the outcome of a random choice between the 6 orderings (which do not change with the number of features). We might expect a random choice to be correct one sixth of the time but it is higher as for some problems there were multiple variable orderings with equally fast timings. The distribution of the computation times: the differences between the computation time of each method and the minimum computation time, given as a percentage of the minimum time, are depicted in Figure 1.

5.1 Range of possible outcomes

The minimum total computing time, achieved if we select an optimal ordering for every problem, is 8 623s (the virtual best heuristic). Choosing at random would take 30 235s, almost 4 times as much. The maximum time, if we selected the worst ordering for every problem, is 64 534s. The K-Nearest Neighbours model (with 78 features) achieved the shortest time of our models and heuristics, with 9 178s, only 6% more than the minimal possible.

Table 4: The top table compares performance of the machine learning classifiers (DT, KNN, MLP, SVM) with the original set of features as published in [24] and the extended set developed in the present paper. For comparison, the bottom table shows the best and worst possible choices, along with the results of a random choice and those of two human-made heuristics (**Brown** and **sotd**).

		DT	KNN	MLP	SVM
From [24] (11 Features)	Accuracy	62.6%	63.3%	61.6%	58.8%
	Time (s)	9 994	10 105	9 822	10 725
New Experiment (78 Features)	Accuracy	65.2%	66.3%	67%	65%
	Time (s)	9 603	9 178	9 399	9 487

	Virtual Best Heuristic	Virtual Worst Heuristic	random	Brown	sotd
Accuracy	100%	0%	22.7%	51%	49.5%
Time (s)	8 623	64 534	30 235	10 951	11 938

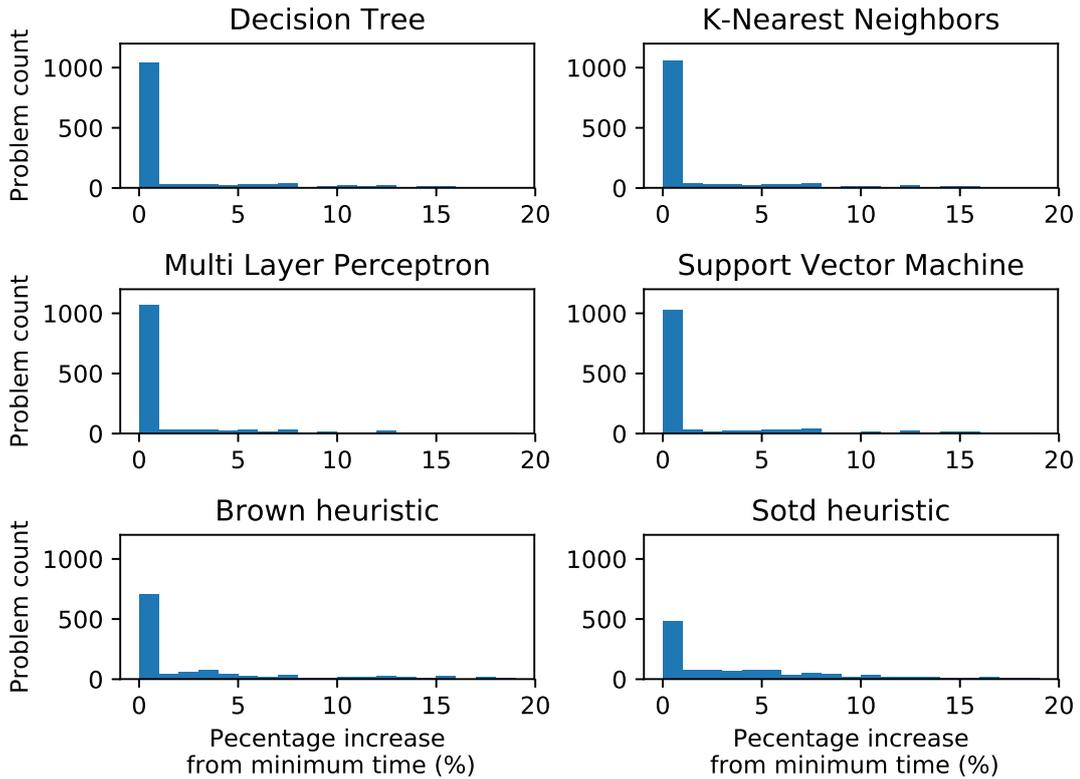


Figure 1: The histograms of the percentage increase in computation time relative to the minimum computation time for each method, calculated for a bin size of 1%.

5.2 Human-made heuristics

Since they are not affected by the new feature framework of the present paper the findings on the human made heuristics are the same as in [24]. Of the two human-made heuristics, **Brown** performed the best, surprising since the **sotd** heuristic has access to additional information (not just the input polynomials but also their projections). Obtaining an ordering for a problem instance with **sotd** hence takes longer than for **Brown** or any ML model – generating an ordering for all problems in the testing dataset took over 30min. Using **Brown** we can solve all problems in 10,951s, 27% more than the minimum. While **sotd** is only 0.7% less accurate than **Brown** in identifying the best ordering, it is much slower at 11 938s or 38% more than the minimum. So, while **Brown** is not much better at identifying the best, it is much better at discarding the worst!

5.3 ML choices

The results show that all ML approaches outperform the human constructed heuristics in terms of both accuracy and timings. Moreover, the results show that the new algorithm for generating features leads to a clear improvement in ML performance compared to using only a small number of human generated features in [24]. For all four modules both accuracy has increased and computation time decreased. The best achieved time was 14% above the minimum using the original 11 features but now only 6% above with the new features.

The computing time for all the methods lies between the best (8 623s) and the worst (64 534s). Therefore, if we scale this time to [0, 100] so that the shortest time corresponds to 0 and the slowest to 100, then the best human-made heuristic (**Brown**) lies at 4.16, and the best ML method (KNN) lies at 0.99. So using ML allows us to be 4 times closer to the minimum possible computing time.

Figure 1 shows that the human-made heuristics result in computing times that are often significantly larger than 1% of the corresponding minimum time for each problem. The ML methods, on the other hand, all result in over 1000 problems ($\sim 75\%$ of the testing dataset) within 1% of the minimum time.

6 Final Thoughts

In this experiment the MLP and KNN models offered the best performance, and a clear advance on the prior state of the art. But we acknowledge that there is much more to do and emphasise that these are only the initial findings of the project and we need to see if the findings are replicated. Planned extensions include: expanding the dataset to problems with more variables and quantifier structure; trying different feature selection techniques, and seeing if classifiers trained for the MAPLE CAD may be applied to other implementations.

Our main result is that a great many more features can be obtained trivially from the input (i.e. without any projection operations) than previously thought, and that these are relevant and lead to better ML choices. Some of these are easy to express in natural language, such as the number of polynomials containing a certain variable, but others do not have an obvious interpretation. This is important because something that is hard to describe in natural language is unlikely to be suggested by a human as a feature, which illustrates the benefit of our framework. This contribution to feature extraction for algebraic problems should be more widely applicable than the CAD variable ordering decision.

Acknowledgements

This work is supported by EPSRC Project EP/R019622/1: *Embedding Machine Learning within Quantifier Elimination Procedures*.

References

- [1] Ábrahám, E., Abbott, J., Becker, B., Bigatti, A., Brain, M., Buchberger, B., Cimatti, A., Davenport, J., England, M., Fontaine, P., Forrest, S., Griggio, A., Kroening, D., Seiler, W., Sturm, T.: SC^2 : Satisfiability checking meets symbolic computation. In: Proc. CISM '16, LNCS 9791, pp. 28–43. Springer (2016), https://doi.org/10.1007/978-3-319-42547-4_3
- [2] Bishop, C.: Pattern Recognition and Machine Learning. Springer (2006)
- [3] Bradford, R., Chen, C., Davenport, J., England, M., Moreno Maza, M., Wilson, D.: Truth table invariant cylindrical algebraic decomposition by regular chains. In: Proc. CASC '14, LNCS 8660, pp. 44–58. Springer (2014), http://dx.doi.org/10.1007/978-3-319-10515-4_4
- [4] Bradford, R., Davenport, J., England, M., Errami, H., Gerdt, V., Grigoriev, D., Hoyt, C., Košta, M., Radulescu, O., Sturm, T., Weber, A.: A case study on the parametric occurrence of multiple steady states. In: Proc. ISSAC '17, pp. 45–52. ACM (2017), <https://doi.org/10.1145/3087604.3087622>
- [5] Bradford, R., Davenport, J., England, M., McCallum, S., Wilson, D.: Truth table invariant cylindrical algebraic decomposition. *J. Symb. Comp.* **76**, 1–35 (2016), <http://dx.doi.org/10.1016/j.jsc.2015.11.002>
- [6] Bradford, R., Davenport, J., England, M., Wilson, D.: Optimising problem formulations for cylindrical algebraic decomposition. In: Intelligent Computer Mathematics, LNCS 7961, pp. 19–34. Springer Berlin Heidelberg (2013), http://dx.doi.org/10.1007/978-3-642-39320-4_2
- [7] Bridge, J., Holden, S., Paulson, L.: Machine learning for first-order theorem proving. *Journal of Automated Reasoning* **53**, 141–172 (2014), <https://doi.org/10.1007/s10817-014-9301-5>
- [8] Brown, C.: Tutorial: Cylindrical algebraic decomposition, at ISSAC '04, (2004), <http://www.usna.edu/Users/cs/wcbrown/research/ISSAC04/handout.pdf>
- [9] Brown, C.: Constructing a single open cell in a cylindrical algebraic decomposition. In: Proc. ISSAC '13, pp. 133–140. ACM (2013), <https://doi.org/10.1145/2465506.2465952>
- [10] Brown, C., Davenport, J.: The complexity of quantifier elimination and cylindrical algebraic decomposition. In: Proc. ISSAC '07, pp. 54–60. ACM (2007), <https://doi.org/10.1145/1277548.1277557>
- [11] Carette, J.: Understanding expression simplification. In: Proc. ISSAC '04, pp. 72–79. ACM (2004), <https://doi.org/10.1145/1005285.1005298>
- [12] Caviness, B., Johnson, J.: Quantifier Elimination and Cylindrical Algebraic Decomposition. Texts & Monographs in Symbolic Computation, Springer-Verlag (1998), <https://doi.org/10.1007/978-3-7091-9459-1>

- [13] Chen, C., Moreno Maza, M.: An incremental algorithm for computing cylindrical algebraic decompositions. In: *Computer Mathematics*, pp. 199–221. Springer Berlin Heidelberg (2014), https://doi.org/10.1007/978-3-662-43799-5_17
- [14] Chen, C., Moreno Maza, M., Xia, B., Yang, L.: Computing cylindrical algebraic decomposition via triangular decomposition. In: *Proc. ISSAC '09*, 95–102. ACM (2009), <https://doi.org/10.1145/1576702.1576718>
- [15] Collins, G.: Quantifier elimination for real closed fields by cylindrical algebraic decomposition. In: *Proc. 2nd GI Conference on Automata Theory and Formal Languages*. pp. 134–183. Springer-Verlag (reprinted in the collection [12]) (1975), https://doi.org/10.1007/3-540-07407-4_17
- [16] Collins, G., Hong, H.: Partial cylindrical algebraic decomposition for quantifier elimination. *Journal of Symbolic Computation* **12**, 299–328 (1991), [https://doi.org/10.1016/S0747-7171\(08\)80152-6](https://doi.org/10.1016/S0747-7171(08)80152-6)
- [17] Davenport, J., Bradford, R., England, M., Wilson, D.: Program verification in the presence of complex numbers, functions with branch cuts etc. In: *Proc. SYNASC '12*, pp. 83–88. IEEE (2012), <http://dx.doi.org/10.1109/SYNASC.2012.68>
- [18] Dolzmann, A., Seidl, A., Sturm, T.: Efficient projection orders for CAD. In: *Proc. ISSAC '04*, pp. 111–118. ACM (2004), <https://doi.org/10.1145/1005285.1005303>
- [19] England, M.: Machine learning for mathematical software. In: *Mathematical Software – Proc. ICMS '18*. LNCS 10931, pp. 165–174. Springer (2018), https://doi.org/10.1007/978-3-319-96418-8_20
- [20] England, M., Bradford, R., Davenport, J.: Improving the use of equational constraints in cylindrical algebraic decomposition. In: *Proc. ISSAC '15*, pp. 165–172. ACM (2015), <http://dx.doi.org/10.1145/2755996.2756678>
- [21] England, M., Bradford, R., Davenport, J., Wilson, D.: Choosing a variable ordering for truth-table invariant CAD by incremental triangular decomposition. In: *Proc. ICMS '14*, LNCS 8592, pp. 450–457. Springer (2014), http://dx.doi.org/10.1007/978-3-662-44199-2_68
- [22] England, M., Davenport, J.: Experience with heuristics, benchmarks & standards for cylindrical algebraic decomposition. In: *Proc. SC² '16*. CEUR-WS **1804** (2016), <http://ceur-ws.org/Vol-1804/>
- [23] England, M., Errami, H., Grigoriev, D., Radulescu, O., Sturm, T., Weber, A.: Symbolic versus numerical computation and visualization of parameter regions for multistationarity of biological networks. In: *Computer Algebra in Scientific Computing (CASC '17)*, LNCS 10490, pp. 93–108. Springer (2017), https://doi.org/10.1007/978-3-319-66320-3_8
- [24] England, M., Florescu, D.: Comparing machine learning models to choose the variable ordering for cylindrical algebraic decomposition. To appear in *Proc. CICM '19* (Springer LNCS) (2019). Preprint: <https://arxiv.org/abs/1904.11061>
- [25] England, M., Wilson, D., Bradford, R., Davenport, J.: Using the Regular Chains Library to build cylindrical algebraic decompositions by projecting and lifting. In: *Mathematical Software – ICMS '14*. LNCS 8592, pp. 458–465. Springer (2014), http://dx.doi.org/10.1007/978-3-662-44199-2_69
- [26] Huang, Z., England, M., Davenport, J., Paulson, L.: Using machine learning to decide when to precondition cylindrical algebraic decomposition with Groebner bases. In: *Proc. SYNASC '16*. pp. 45–52. IEEE (2016), <https://doi.org/10.1109/SYNASC.2016.020>
- [27] Huang, Z., England, M., Wilson, D., Bridge, J., Davenport, J., Paulson, L.: Using machine learning to improve cylindrical algebraic decomposition. *Mathematics in Computer Science*, Volume to be assigned, 28 pages (2019), <https://doi.org/10.1007/s11786-019-00394-8>
- [28] Huang, Z., England, M., Wilson, D., Davenport, J., Paulson, L., Bridge, J.: Applying machine learning to the problem of choosing a heuristic to select the variable ordering for cylindrical algebraic decomposition. In: *Intelligent Computer Mathematics*, LNAI 8543, pp. 92–107. Springer International (2014), http://dx.doi.org/10.1007/978-3-319-08434-3_8

- [29] Iwane, H., Yanami, H., Anai, H., Yokoyama, K.: An effective implementation of a symbolic-numeric cylindrical algebraic decomposition for quantifier elimination. In: Proc. SNC '09, pp. 55–64. SNC '09 (2009), <https://doi.org/10.1145/1577190.1577203>
- [30] Jovanovic, D., de Moura, L.: Solving non-linear arithmetic. In: Gramlich, B., Miller, D., Sattler, U. (eds.) Automated Reasoning – Proc. IJCAR '12, LNCS 7364, pp. 339–354. Springer (2012), https://doi.org/10.1007/978-3-642-31365-3_27
- [31] Kobayashi, M., Iwane, H., Matsuzaki, T., Anai, H.: Efficient subformula orders for real quantifier elimination of non-prenex formulas. In: Proc. MACIS '15, LNCS 9582, pp. 236–251. Springer International Publishing (2016), https://doi.org/10.1007/978-3-319-32859-1_21
- [32] Kühlwein, D., Blanchette, J., Kaliszyk, C., Urban, J.: MaSh: Machine learning for sledgehammer. In: Interactive Theorem Proving, LNCS 7998, pp. 35–50. Springer Berlin Heidelberg (2013), https://doi.org/10.1007/978-3-642-39634-2_6
- [33] Kuipers, J., Ueda, T., and Vermaseren, J.A.M.: Code optimization in FORM. Computer Physics Communications, **189**, 1–19 (2015), <https://doi.org/10.1016/j.cpc.2014.08.008>
- [34] Liang, J., Hari Govind, V., Poupart, P., Czarnecki, K., Ganesh, V.: An empirical study of branching heuristics through the lens of global learning rate. In: Proc. SAT '17, LNCS 10491, pp. 119–135. Springer (2017), https://doi.org/10.1007/978-3-319-66263-3_8
- [35] McCallum, S.: An improved projection operation for cylindrical algebraic decomposition. In: [12], pp. 242–268. (1998), https://doi.org/10.1007/978-3-7091-9459-1_12
- [36] Markowski, C.A. and Markowski, E.P.: Conditions for the effectiveness of a preliminary test of variance. The American Statistician, **44**:4, 322–326 (1990)
- [37] McCallum, S., Parusiński, A., Paunescu, L.: Validity proof of Lazard’s method for CAD construction. Journal of Symbolic Computation **92**, 52–69 (2019), <https://doi.org/10.1016/j.jsc.2017.12.002>
- [38] Mulligan, C., Bradford, R., Davenport, J., England, M., Tonks, Z.: Non-linear real arithmetic benchmarks derived from automated reasoning in economics. In: Proc. SC² '18, pp. 48–60. CEUR-WS **2189** (2018), <http://ceur-ws.org/Vol-2189/>
- [39] Mulligan, C., Davenport, J., England, M.: TheoryGuru: A Mathematica package to apply quantifier elimination technology to economics. In: Mathematical Software – Proc. ICMS '18. LNCS 10931, pp. 369–378. Springer (2018), https://doi.org/10.1007/978-3-319-96418-8_44
- [40] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. Journal of Machine Learning Research **12**, 2825–2830 (2011), <http://www.jmlr.org/papers/v12/pedregosa11a.html>
- [41] Strzeboński, A.: Cylindrical algebraic decomposition using validated numerics. Journal of Symbolic Computation **41**(9), 1021–1038 (2006), <https://doi.org/10.1016/j.jsc.2006.06.004>
- [42] Sturm, T.: New domains for applied quantifier elimination. In: Computer Algebra in Scientific Computing, LNCS 4194, pp. 295–301. Springer (2006), https://doi.org/10.1007/11870814_25
- [43] Urban, J.: MaLAREa: A metasystem for automated reasoning in large theories. In: Proc. ESARLT '07, CEUR-WS **257**, p. 14. CEUR-WS (2007), <http://ceur-ws.org/Vol-257/>
- [44] Wilson, D., Bradford, R., Davenport, J., England, M.: Cylindrical algebraic sub-decompositions. Mathematics in Computer Science **8**, 263–288 (2014), <http://dx.doi.org/10.1007/s11786-014-0191-z>
- [45] Wilson, D., Davenport, J., England, M., Bradford, R.: A “piano movers” problem reformulated. In: Proc. SYNASC '13, pp. 53–60. IEEE (2013), <http://dx.doi.org/10.1109/SYNASC.2013.14>
- [46] Xu, L., Hutter, F., Hoos, H., Leyton-Brown, K.: SATzilla: Portfolio-based algorithm selection for SAT. J. Artificial Intelligence Research **32**, 565–606 (2008), <https://doi.org/10.1613/jair.2490>