Explainable AI Insights for Symbolic Computation: A case study on selecting the variable ordering for cylindrical algebraic decomposition

Lynn Pickering[a], Tereso Del Rio Almajano[b], Matthew England[b], Kelly Cohen[a]

[a]*Aerospace Engineering and Engineering Mechanics*
University of Cincinnati, Cincinnati, 45221, Ohio, USA
[b]*Research Centre for Computational Science and Mathematical Modelling*
Coventry University, Coventry, CV1 5FB, West Midlands, UK

Abstract

In recent years there has been increased use of machine learning (ML) techniques within mathematics, including symbolic computation where it may be applied safely to optimise or select algorithms. This paper explores whether using explainable AI (XAI) techniques on such ML models can offer new insight for symbolic computation, inspiring new implementations within computer algebra systems that do not directly call upon AI tools. We present a case study on the use of ML to select the variable ordering for cylindrical algebraic decomposition. It has already been demonstrated that ML can make the choice well, but here we show how the SHAP tool for explainability can be used to inform new heuristics of a size and complexity similar to those human-designed heuristics currently commonly used in symbolic computation.

Keywords: Explainable AI, Computer Algebra, Heuristic Development, Cylindrical Algebraic Decomposition, Variable Ordering
2000 MSC: 68W30, 68T20

1. Introduction

1.1. Machine learning and mathematics

Machine Learning (ML) refers to statistical techniques that learn rules from data. They allow computer systems to improve their performance on a task without any change to their explicit programming. ML underpins the prominent AI advances of recent years, driven forward by the growth in available data and computing power, new ML architectures, and hardware designed specifically to them.

There are increasing attempts to use ML in mathematics. For example, Lample and Charton [37] trained a transformer to integrate expressions and find analytical solutions to differential equations; while in a survey article He [28] details a variety of attempts to use ML to predict properties from mathematical structures such as groups and graphs. We note a recent special issue of the Journal of Symbolic Computation dedicated to *Algebraic Geometry and Machine Learning* [27] including ML approaches to e.g. find the real discriminant locus [2]. However, it has been observed that mathematical and logical reasoning is an area of particular difficulty for ML, especially the recent natural language tools such as ChatGPT[1].

Symbolic computation is not an obvious domain in which to apply ML. The Computer Algebra Systems (CASs) which implement symbolic computation algorithms are mainly designed to produce exact answers; rarely using numeric computation in preference for symbolic solutions. Lample and Charton [37] reported that their Transformer[2] was able to integrate far more examples correctly within a time limit than various CASs. However, this analysis combined two very different cases of failure: when a solution could not be found within a time limit and when the wrong solution is found. While the CASs may not solve problems quickly, they should *never* produce the wrong answer. Thus there is little appetite amongst CAS developers to use ML for directly producing output.

---

[1]See e.g. this blog post by S. Wolfram: https://writings.stephenwolfram.com/2023/01/wolframalpha-as-the-way-to-bring-computational-knowledge-superpowers-to-chatgpt/

[2]A particular design of ML model that has found great success in natural language processing tasks: it is used for example by ChatGPT.

1.2. Safe use of machine learning in computer algebra systems

This does not mean that ML has nothing to offer Symbolic Computation. CAS algorithms often come with choices that have no effect on the mathematical correctness of the end result but can have a big impact on the resources required. Further, they can affect how the end result is presented (i.e. two very different, but mathematically equivalent, expressions). Consider for example the algorithm of Buchberger [9] to produce the Gröbner Basis for an ideal. It does not specify the order in which S-pairs are studied, the order in which the corresponding S-polynomial is reduced by the generating set, the monomial ordering to be used, nor the underlying variable ordering. Any decision for these choices allows the production of a correct Gröbner Basis but each decision affects the size of the basis produced, the polynomials within, and the time taken to compute them.

Sometimes there are well-documented strategies on how CASs make such choices (e.g. sugar selection for S-pairs [25]) but in many cases, they rely on human-made heuristics or "*magic constants*" [10] which are often not scientifically validated or even documented. Such choices are good candidates for ML: the underlying relationships are complex but not the main object of study. The earliest examples of ML for CAS optimisation known to the authors are: in 2014 Huang et al. [31] used a support vector machine to choose the variable ordering for cylindrical algebraic decomposition; in 2015 Kuipers et al. [36] used a Monte-Carlo tree search to find the representation of polynomials that are most efficient to evaluate; and in 2016 Simpson et al. [53] used ML classifiers to pick from various algorithms that compute the resultant, Huang et al. [30] used ML to decide whether to precondition input for cylindrical algebraic decomposition, and Kobayashi et al. [34] used ML to decide the order of sub-formulae solving for real quantifier elimination.

In all of these examples, the ML solution outperformed the previous human-made heuristics for the choices, at least on the dataset used. The algorithms of symbolic computation are often exponential in their worst complexity but the average case complexity is less widely studied. It is possible that these ML optimisations may offer a route to avoid the worst case, at least for many applied examples. Thus there is great value in their continued study and exposition. However, the motivation of the present paper is whether ML can offer anything more than these efficiency gains. A recent article in Nature [16] suggested that ML can help pure mathematicians with the development of new theorems. Can ML offer such insight to the developers of symbolic computation algorithms?

1.3. Explainable artificial intelligence

Explainable AI (XAI) refers to ML techniques which offer an explanation as to how an AI decision was made. It has become a very active area of research, see for example DARPA's XAI programme [26]. The role of XAI is usually to allow for more effective use of the AI tools, and the building of user trust in the tools. But in our case, we hypothesise that XAI could reveal workings of the ML model that are useful to inform further development of the underlying symbolic computation.

We were motivated by the work of Peifer et al. [48] developing an ML solution to the problem of selecting S-pairs in Buchberger's algorithm. They trained an "*agent*" (ML model) to make the decision using reinforcement learning[3] based on the number of polynomial additions required to process the choice. This agent outperformed existing strategies such as [25]. We draw attention to [48, Section 5.1] which tried to explain the agents' strategy (albeit without employing XAI tools). It identified some simple components such as a preference for pairs whose S-polynomials are monomials and a preference for pairs whose S-polynomials are low degree. Firstly, these components suggested decisions made on the basis of the S-polynomials rather than the S-pairs [4] which are used by all the human-made heuristics. Thus this opens a whole new category of heuristics to consider, perhaps leading to a greater understanding of the problem. Secondly, experiments using these strategies alone showed they outperformed the existing heuristics (but not the full ML agent). Thus the analysis suggests new "*human-level*" heuristics, i.e. heuristics of a similar size to those that are human-made and that, once discovered, can be implemented without any AI tools.

The present paper seeks to identify if XAI tools can offer an automated way to make such analyses, to gain new understanding of these choices in symbolic computation algorithms.

1.4. Contributions and plan of the paper

In this paper, we return to the problem of selecting the variable ordering for cylindrical algebraic decomposition, the first CAS optimisation by ML that was considered in [31]. We introduce the necessary background information on this problem in Section 2. We apply the popular SHAP methodology

---

[3]With reinforcement learning, instead of having a labelled dataset to train against, the agent repeatedly makes decisions which receive a score that they seek to optimise.

[4]even though it was properties of the pairs fed to the network.

for XAI to an existing ML pipeline developed for the problem. We introduce SHAP in Section 3 and report on the results of its application on this problem in Section 4. In Section 5 we process the SHAP data to suggest some XAI-recommended human-level heuristics for the problem and finally, in Section 6 we report on the performance of these.

We find that one of these is equal to the current state-of-the-art human-developed heuristic for the problem [50] and that a human-level combination of these XAI selected features allows us to outperform prior work. Thus our contributions are (a) a new state-of-the-art human-level heuristic and perhaps more importantly; (b) evidence for a new methodology to employ in CAS development that uses AI in algorithm design but not implementation.

2. Cylindrical Algebraic Decomposition

2.1. Cylindrical algebraic decomposition

Cylindrical Algebraic Decomposition (CAD) is an important symbolic computation algorithm, proposed by Collins [13] in 1975. Given a set of multivariate polynomials and a variable ordering, CAD will decompose the corresponding real space into connected regions, called cells, such that within each the sign of each polynomial is invariant (i.e. positive, negative or zero). Each cell is semi-algebraic, meaning it may be defined by a set of polynomial constraints; and the cells are arranged in cylinders with respect to the variable ordering, meaning the projections of any two cells onto a lower dimensional space with respect to the ordering are either equal or disjoint.

CAD has found applications in various fields, ranging from robotics [41] through economics [45] to biology [51]. Its most commonly used for, and was proposed by Collins [13] as, a sub-method to perform real Quantifier Elimination (QE): i.e. to transform a formula in first-order logic whose atoms are Boolean constraints to an equivalent one without any quantifiers. The decomposition into cells allows for studying the problem at a finite number of points, while the cylindrical structure allows for easy projection and negation of the cells to construct the quantifier-free solution.

However, CAD has a worst-case theoretical complexity that grows doubly exponential in the number of variables [15]; something regularly encountered in practice and thus vastly reducing the scope of use cases for the algorithm. There have thus been many lines of research to improve CAD, such as improved projection operators [43], partial CAD [14], equational constraints
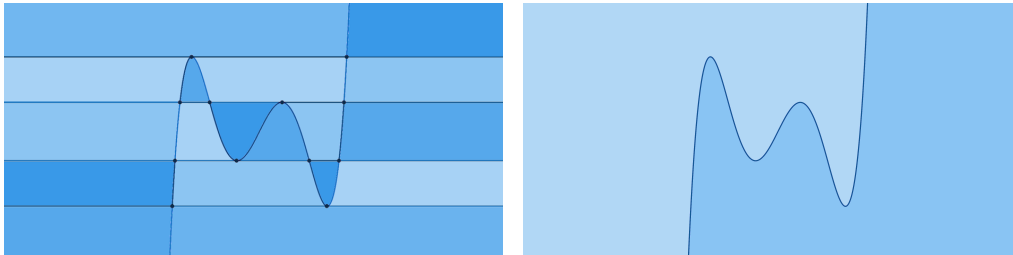
Figure 1: CADs sign-invariant for the set of polynomials $\{x^5 + 5x^4 + 5x^3 - 5x^2 - 6x - 2y\}$. Using ordering $x \succ y$, we obtain a CAD with 57 cells (the 18 shaded areas, 27 lines and 12 points). Using the ordering $y \succ x$ generates only 3 cells (2 areas and one line).

[19], symbolic-numeric lifting [54], alternative computation schemes via triangular decomposition [11] and incremental processing of polynomials [35], and entirely new repackaging of the CAD theory [7], [33].

2.2. CAD variable ordering

CAD, including all the optimised and extended forms listed above, requires there to be a declared ordering on the variables. This controls both the order of operations in the algorithms and the form of the output (being used in the definition of the cylindrical structure). Depending on the application, there may be a free or constrained choice for the ordering. For example, when using CAD for real QE the variables must be ordered as they are quantified: but variables in blocks of the same quantifier (and free variables) can be swapped with each other.

This choice of variable ordering has been shown to have a huge impact, both practically [17] and in terms of theoretical complexity: Brown and Davenport [8] showed there exists a family of polynomial sets such that the complexity of CAD can be constant or doubly exponential depending on the chosen variable ordering. It is possible to appreciate the importance of variable ordering in the simple two-variable example of Figure 1, where one of the orderings generates three times more cells than the other.

2.3. Prior human-designed heuristics for choosing the ordering

Once the importance of variable ordering was established, researchers began looking for strategies to choose a good variable ordering. The first attempts consisted of designing human-made heuristics. In 2004 Brown [6] documented a heuristic based on three simple hand-picked features of the set of polynomials, for the software QEPCAD. We call this heuristic `Brown`.

6

There have been other heuristics produced which can offer greater accuracy but at greater expense, performing increasing numbers of steps in the CAD algorithm: Dolzmann et al. [17] concluded it best to perform the projection stage of CAD and compare sums of the total degree of the polynomials produced (`sotd`); Bradford et al. [5] considered the initial decomposition of the real line; and Wilson et al. [55] the open cells in the decomposition. It is perhaps surprising that the original heuristic of Brown is able to perform competitively or even beat these others despite having access to less information [32]. This suggests the necessary information to identify a good ordering may be available from the input alone.

Most recently del Río and England [50] designed a new simple heuristic `gmods` based on properties of the input polynomials, selected by studying which features of the input polynomials have the greatest impact on the complexity analysis of CAD.

2.4. Prior AI-designed heuristics for choosing the ordering

In the last decade machine learning models have also been trained to select the variable ordering for CAD. The first attempt was made by Huang et al. [31] in 2014 who used a support vector machine for the task. Later England and Florescu [20] experimented with a wider range of models, methods for feature engineering [21] and improved metrics for training [22], culminating in a machine learning pipeline available to use for the task [23]. Separately Chen et al. [12] experimented with deep learning for variable selection.

In all cases, these ML models outperform existing heuristics. However, CAS developers may still be reluctant to include a method trained using ML in their software, possibly because they create a dependence on external code, or more likely due to concerns about whether over-fitting on the training data will lead to poor generalisation. Another issue is the lack of explanation these AI models are able to give about their choices.

We will experiment later using the tools for an AI selection of the CAD variable ordering presented in [23]. In that paper, the authors present a pipeline to train four ML models on the task. When specified to three variables polynomial sets this utilises the ideas described in [21] to represent such polynomials with 81 features (which evaluate to floating point values). It is these numerical feature vectors that are presented to the ML models.

In this paper, we propose an approach that is somewhere in between those of this subsection and the previous one. Rather than handpicking

some features as in Section 2.3, will we start with the algorithmically produced features of Section 2.4 but then use XAI techniques to extract the most important features from the AI and use these to design heuristics that could easily be understood by humans and embedded into any CAD implementation without use of AI software.

3. Explaining Machine Learning Models

XAI is often motivated as a means to build trust in the tool among users. But there are certainly other reasons we may require an explanation of an ML model. Perhaps the model is making a poor decision and the data scientist wishes to understand why so as to improve it; or perhaps the explanation tells us more about the underlying trends of data leading to new research directions. The latter is the motivation for explainability in this work. Within the XAI field there is a growing emphasis on designing models that are intrinsically explainable[5], however, the most common approach in use today is post-hoc explainability whereby a model is analysed after its creation and an explanation is generated. The review article by Du et al. [18] offers a good overview of XAI.

Since we work with the existing ML pipeline [23] already designed for the problem we do not consider editing the ML itself for explainability but instead revert to the more common post-hoc analysis. That pipeline works with several ML models, and so we desire a model-agnostic method for explainability. The following is the current state-of-the-art process for this.

3.1. SHAP

Molnar [44] gives an overview of many explainability methods along with their advantages and disadvantages. SHAP (SHapley Additive exPlanations), a model agnostic method, [40] is chosen as our method of post-explainability analysis for identifying the features of most use to the ML models. The SHAP approach was introduced in 2017 Lundberg and Lee [40] as a unification of six existing explanation methods already present in the literature: LIME, DeepLIFT, Layer-Wise Relevance Propagation, classic

---

[5]There are well-established views on which models are intrinsically explainable and an assumed tradeoff against performance with e.g. neural networks offering the best performance but the least explanation. However, in practice this trade-off is not always so clear, see e.g. [29].

Shapley regression values, Shapley sampling values, and Quantitative Input Influence. Because SHAP unifies these various methods, it can be used for a global and local explanation of models: i.e. it can offer explanations for performance on an individual problem or an entire dataset.

The name and theory behind SHAP is based on the classical concept of Shapley values from game theory, which quantifies the contribution of each player to a game [52]. It makes use of an explainer model, which is an approximation of the original model. The idea is to figure out the contribution of each player to the score of the game by playing the game with all the combinations of players, or the power set of the players. Take an example game with three players: player 1, player 2, and player 3. Each player plays the game by themselves, then each player plays the game with each of the other two, and finally they all play the game together. We calculate the difference between the final scores of each game where player 1 was not present, and when player 1 was present under a weighting to get the contribution of player 1 to the game. In our application of SHAP, the features of our polynomial sets are the players and the ML models are the games. SHAP removes and adds each feature to the model to see what effect that feature has.

Since we are working over the power set of all the features in the model, such a calculation becomes very expensive quickly: exponential on the number of features in a model [40]. However, the SHAP open source package [39] uses many estimators and powerful shortcuts to make the extensive power set of calculations possible for larger models. For example, Kernel SHAP, used to analyse all the models other than decision tree in this work, uses a special weighted linear regression in computing a feature's importance, as outlined in the README of the GitHub page for SHAP. For the decision tree model, the tree explainer is used. This is a fast method to estimate SHAP values for tree models.

3.2. SHAP waterfall plots

Figure 2 shows a SHAP waterfall plot, which is one visualization in SHAP for understanding a model's decision on a single instance. The ML model studied here was to predict if a person's income is less than or greater than $50k annually [38]. Negative values imply probabilities of less than 0.5 that a person has income over $50k. In this case, the features that contribute to a prediction that a person makes less than $50k are the blue bars, while the features that contribute to a prediction that a person makes more than $50k are the red bars. The features from top to bottom are those that have the
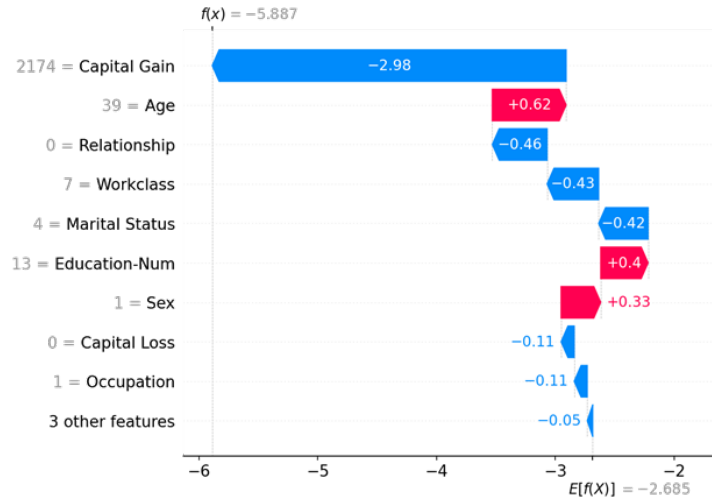
Figure 2: SHAP waterfall plot taken from [38]: explains a single instance of an XGBoost model trained on the classic UCI adult income dataset.

largest effect on the final model output. In the waterfall plots, $E[f(X)]$ is the base value and the number that each decision starts at, because it is the average of all the output across all the instances, which would be the best guess a model would have for this dataset if there were no features. The grey values next to each feature are the values of the corresponding feature for that instance.

3.3. Multi-class SHAP

The ML pipeline we experiment with is designed for multi-classification, i.e. classification with more than two output classes. SHAP may also be applied for multi-classification. It simply creates a binary model for each output, with the binary SHAP model that reaches an output closest to one for an instance giving the prediction.

4. SHAP Applied to ML for CAD Variable Ordering

We work with the feature engineering approach to the problem developed by Florescu and England [21]. This algorithmically generates features to represent systems of polynomials to ML models. It was inspired by the earlier hand-picked features based on variable degree and sparsity [6], [31]

10

and generalised this to the widest set of such features that could be generated in a similar way. When applied to three variable polynomial systems this framework creates 81 distinct features, which the pipeline in [23] uses as the input to train four ML models (implemented in Python from the scikit-learn package [47]) which are: Decision Tree (DT), K-Nearest Neighbors (KNN), Support Vector Machine (SVM), and Multi-Layer Perceptron (MLP)[6]. The output that each model predicts is a class (1−6) representing one of the six possible variable orderings of three variables. In these papers, the experiments showed the ML models to all outperform the best performing human-made heuristic of the time, that of Brown [6].

The work described in this section uses SHAP to gain insight into how those ML models from [23] achieve this performance. We seek to understand whether the ML models find the most important features to be those selected by humans, which are then supplemented with additional information; or whether they instead rely on some previously overlooked feature. Any prominent features previously unconsidered will be studied later in the paper. This section first looks at the decision-making process for a single test instance to gain a local understanding of a particular ML model (the MLP). The process from the local decision to the global decision is then explained, and the top five features across the entire model, for each model are identified and discussed.

4.1. Methodology

The methodology used in applying SHAP to ML algorithms for picking the CAD variable ordering is described next, starting with the data set used to train the algorithms, the specific manner in which SHAP was implemented, and details on the features used as input to the ML algorithms.

4.1.1. Dataset

We started by recreating the experiments of Florescu and England [23]. These used a dataset of random problems generated with the same statistical properties as the industrial problems hosted in the QF_NRA category of the SMT-LIB [1]. We identified that the dataset used in Florescu and England [23] had a majority of problems with class 5 or 6 as the output class. In this paper, we applied the additional step of balancing the dataset so that

---

[6]We refer a reader seeking the details of these models to e.g. the textbook Bishop [3].

the classes are more equal in size. We take advantage of the fact that the nature of a polynomial set does not change with the names of its variables. This allows for the random permutation of variable names (x_1, x_2, x_3), and a corresponding change of the best ordering for the output, to obtain a balanced version of the dataset. As a reference of which output class number corresponds to which order of the variables, we refer to Table 1.

Table 1: The six possible variable orderings for our dataset

Ordering Name	Ordering
Ordering 1	$x_1 \succ x_2 \succ x_3$
Ordering 2	$x_1 \succ x_3 \succ x_2$
Ordering 3	$x_2 \succ x_1 \succ x_3$
Ordering 4	$x_2 \succ x_3 \succ x_1$
Ordering 5	$x_3 \succ x_1 \succ x_2$
Ordering 6	$x_3 \succ x_2 \succ x_1$

We note that while this renaming makes no difference to the meaning of the dataset, training an ML model on an unbalanced dataset can lead to the ML model simply choosing the class that occurs the most often, and being rewarded for higher accuracy in this way. In general, we might expect this to lead to overfitting, where the model may perform far worse on data from a different source. In our experiments the Decision Tree was found to have particularly suffered from this, making far poorer choices when trained with the balanced dataset, however, the other three models performed similarly well with balanced and unbalanced data and in particular, still outperformed the human-designed heuristics in the unbalanced case.

From the point of view of this paper, balancing is particularly important for explainability. Our initial SHAP analysis on the models trained on the unbalanced dataset clearly reflected a simplistic preference for the majority classes and gave us little insight into understanding the important features of this application.

Table 2: Wall clock runtimes of the SHAP analysis for each model with 500 training points and 100 test points. CPU used is Intel(R) Xeon(R) CPU E5-1620 v3 @ 3.50GHz

Model	Total Run Time (HH:MM:SS)
KNN	07:02:42
MLP	00:09:52
DT	00:00:10
SVM	11:38:43

4.1.2. SHAP

SHAP is run on the four ML models. Table 2 shows the very significant differences in SHAP run-time between the models. This is due to both the nature of how long a single ML model takes to train, and the approximations that the SHAP library is able to make use of. We note that SHAP is a procedure we are running once to gain insight, and so are not greatly concerned with its one-off cost. Due to the long-running time for SHAP, 500 training points are used for each SHAP calculation, chosen randomly from the 2947 total training instances used to train the full models described above, and the SHAP calculations are performed on just 100 test points. This gives sufficient insight into the model decisions without requiring more than 12 hours of calculation time. We refer back to Section 3.1 for information on how SHAP makes its calculations.

4.1.3. Features

The derivation of the features used in the dataset is explained in [21]. In brief, instead of handpicking a set of features we algorithmically generate multiple features. This is done by acting upon a list of lists of lists with the innermost list capturing the degrees of variables in a monomial, which are gathered together next in polynomials, and then finally for the whole system. The actions performed on these are simple (cheap) operations such as taking the maximum, sum or average of the numbers. By looking at all possible combinations this method generates 81 distinct features for three variable systems, from which the ML models can choose what they find most useful. We define in Table 3 some of the terminology used in creating these features.

For example, $sum(max(v_1(S)))$ is one of the features that multiple models place importance upon. Its direct translation into plain English is "*the sum across all polynomials of the maximum of the degrees of its monomials for the variable x_1*", which simplifies to, "*the sum across all polynomials of their degrees for the variable x_1*". Similarly, the meaning of $avg(sg(max(v_1(S))))$ is "*the average across all polynomials of the sign of the maximum of the degrees of its monomials for the variable x_1*", which simplifies to "*the average number of polynomials that contain the variable x_1*".

Table 3: Translation of Features

Notation	Text Translation
S	Lists of polynomials e.g. $S = [x_1, x_1^2 - 2x_1x_2^2 + x_2^2 - 3]$
$v_i(S)$	A list that has for each polynomial in S a list of the degrees of x_i in its monomials e.g. $v_2(S) = [[0], [0, 2, 2, 0]]$
$sv_i(S)$	A list that has for each polynomial in S a list of the total degrees in its monomials if x_i appears in the monomial, and 0 otherwise. e.g. $sv_1(S) = [[1], [2, 3, 0, 0]]$
$op(L)$	If L is a list it performs this syntax means to perform operation op upon it. If L is a list of lists it returns a list containing the the result of the operation for each of the original lists. e.g. $max(avg(sv_1(S))) = max([1, 5/2]) = 5/2$
$sg(L)$	This returns the same structure as L where each numerical value is substituted by its sign. e.g. $sum(sg(avg(sg(v_2(S))))) = sum(sg([0, 1/2]))$ $\qquad\qquad\qquad\qquad = sum([0, 1]) = 1$

4.2. SHAP Results

In this section, we interpret the results from SHAP when applied to our problem using the methodology above. First, we explain how it can be used to give an explanation for the classification of a single problem instance: a local SHAP analysis. Then we aggregate across all instances to give an overview of what SHAP tells us about the most relevant features for each model over the dataset: a global SHAP analysis.

4.2.1. Local SHAP Analysis

For each ML method, for each of the six binary classification models, and for each instance, SHAP can produce a waterfall plot (see 3.2) to show the effect each feature has on the final output. As an example, let us look in detail at the MLP model for a particular test instance. An explanation for the binary classifier of whether or not the instance belongs to class 5 (selects Ordering 5) is given via the waterfall plot in Figure 3. The blue arrows in the waterfall indicate that the feature is pushing the prediction towards an output of 0 (that this is the wrong ordering), and the red arrows that the feature pushes towards an output of 1 (that this is the right ordering).

The binary classifier suggested the instance should be classified as Ordering 5 with a probability of 0.71. This was the highest probability among the six binary classifiers and so Ordering 5 was selected as the classification. Orderings $1 - 4$ get a score very close to 0. Ordering 6 receives a score much higher than 0, but not as high as Ordering 5. This is not surprising: Ordering 5 and 6 both select the same variable to project first and it has been observed previously that this has the greatest impact on the calculation time of the CAD.

4.2.2. Global SHAP Analysis

To get the global importance of the features for each model over the dataset we may aggregate the absolute values of the contribution of the features across all instances and binary classifiers. After this, the five most important features for each model found by SHAP are given in Table 4.

Plots of the relative importance of features to each model by class are given in Figure 4. These plots give a good insight into not only that feature's impact but also how a model trains and learns from the data. The KNN model, for example, creates little distinction between the impacts of the features. This is what we would expect from a KNN, because the model is looking for those instances which are close to it to make a decision, and it is
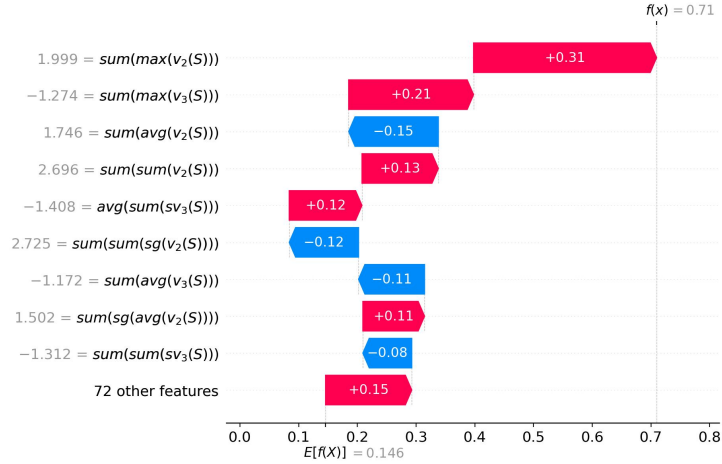
Figure 3: An explanation of a decision made by the MLP model on an example CAD problem instance, for the selected output ordering, ordering 5: $x_3 \succ x_1 \succ x_2$

the aggregate of the number of instances close to a certain test point that will decide the output. This applies similarly to the SVM. In contrast, the MLP and DT models show more distinction in the importance of these features: decision trees in particular are known to place importance on a few features.

Upon close examination of Figure 4, some interesting phenomena may be observed. We see that often the impact on the output class for features of a certain variable is largest if that variable is last in the ordering of that output class. For example, the feature with the greatest impact on the KNN model is an x_3 feature, and Orderings 1 and 3 have the highest proportional impact which (referring back to Table 1) are the orderings with x_3 last. We see this for the x_2 features and the SVM model as well. But the MLP model shows different behavior. Here, only x_2 and x_3 features are seen, and both Ordering 3 and 5 have the greatest impact.

4.2.3. Features in human-designed heuristics

The Brown heuristic [6] uses three simple metrics to choose a variable ordering, all of which may be described in the language of Section 4.1.3:

- $max(max(v_i(S)))$,

- $max(max(sv_i(S)))$, and
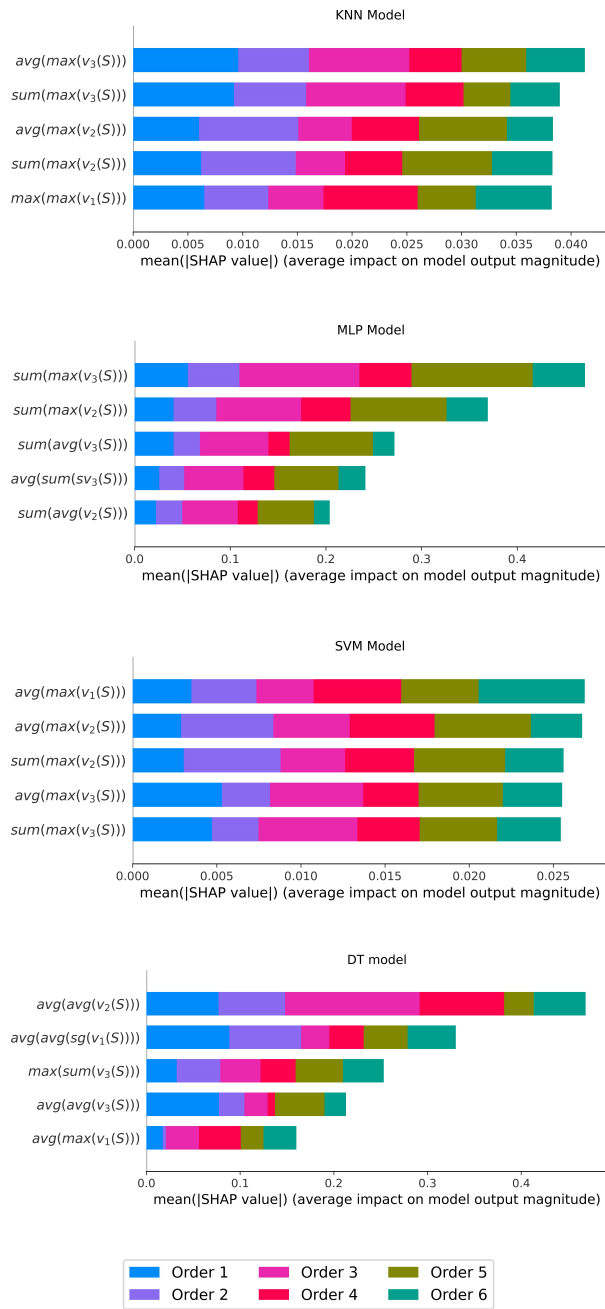
- $sum(sum(sg(v_i(S))))$.

16

Figure 4: The top five features for each model from aggregating over the 100 test points. Ordered by average impact on model output magnitude, and colored by output class.

FI	KNN	MLP
1	$avg(max(v_3(S)))$	$sum(max(v_3(S)))$
2	$sum(max(v_3(S)))$	$sum(max(v_2(S)))$
3	$avg(max(v_2(S)))$	$sum(avg(v_3(S)))$
4	$sum(max(v_2(S)))$	$avg(sum(sv_3(S)))$
5	$max(max(v_1(S)))$	$sum(avg(v_2(S)))$

FI	SVC	DT
1	$avg(max(v_1(S)))$	$avg(avg(v_2(S)))$
2	$avg(max(v_2(S)))$	$avg(avg(sg(v_1(S)))$
3	$sum(max(v_2(S)))$	$max(sum(v_3(S)))$
4	$avg(max(v_3(S)))$	$avg(avg(v_3(S)))$
5	$sum(max(v_3(S)))$	$avg(max(v_1(S)))$

Table 4: The top five features that each model found to have the highest average impact on model output magnitude.

We cannot compare directly to the features of the other human-designed heuristics in Section 2.3 as they all make use of expensive algebraic operations rather than using features from the input polynomials alone.

5. Post SHAP analysis to identify the most relevant features

SHAP produced for us a global analysis explaining the features that are important to each of the four ML models we experiment with. We now work to create a unified ranking of the features across the four models for the purpose of creating new human-level heuristics in Section 6.

5.1. Merging features that generate the same heuristic

We aim for heuristics independent of particular variable names which are present in the ML features. I.e. the prominence of $sum(sum(v_1(S)))$ in ML will lead us to the same heuristic as $sum(sum(v_3(S)))$. So we now aggregate the contribution of the features across the three variables, resulting in rankings like Table 5 for each model. This reduces the original 81 features to 27. Here $v_i(S)$ is a short notation to denote the combination across the three variables, i.e. $sum(sum(v_i(S))) := \sum_{i=1}^{3} sum(sum(v_i(S)))$.

18

Feature Name	Summed SHAP Value
$sum(max(v_i(S))))$	90.716
$sum(avg(v_i(S))))$	57.566
$sum(max(sv_i(S))))$	41.690
$sum(sum(v_i(S))))$	36.670
$avg(avg(sg(v_i(S)))))$	35.676
$avg(sum(sv_i(S))))$	34.758
$avg(max(sv_i(S))))$	30.981
$avg(avg(v_i(S))))$	29.248
$sum(sum(sv_i(S))))$	28.757
$avg(avg(sv_i(S))))$	26.053
$sum(sum(sg(v_i(S)))))$	25.164
$sum(sg(avg(v_i(S)))))$	23.751
$sum(avg(sv_i(S))))$	21.860
$avg(sum(sg(v_i(S)))))$	21.495
$sum(avg(sg(v_i(S)))))$	18.058
$avg(sg(sum(v_i(S)))))$	16.486
$max(max(v_i(S))))$	15.033
$max(avg(v_i(S))))$	13.687
$max(sum(sv_i(S))))$	13.443
$max(max(sv_i(S))))$	12.029
$avg(max(v_i(S))))$	11.864
$max(max(sg(v_i(S)))))$	11.586
$max(avg(sv_i(S))))$	11.559
$max(sum(v_i(S))))$	11.389
$max(sum(sg(v_i(S)))))$	10.451
$avg(sum(v_i(S))))$	9.318
$max(avg(sg(v_i(S)))))$	7.493

Table 5: Features important for the Multi Layer Perceptron after merging across the variables.

Moreover, some features, even though they produce different values for the instances, share some relations. In particular, some of these features are proportional to each other (e.g. $sum(max(v_1(S)))$ is going to be a multiple of $(avg(max(v_1(S))))$ for every set of polynomials and therefore they can be seen theoretically to induce the same heuristics as designed in Section 6. This additional merging reduces the number of features further to 18. The ranking of such features for the MLP model is given in Table 6. Similar rankings are available for the other three models.

Feature Name	Summed SHAP Value
$sum(max(v_i(S))))$	102.58
$avg(avg(v_i(S))))$	86.814
$sum(max(sv_i(S))))$	72.671
$sum(sum(sv_i(S))))$	63.515
$avg(avg(sg(v_i(S)))))$	53.735
$avg(avg(sv_i(S))))$	47.913
$sum(sum(sg(v_i(S)))))$	46.66
$sum(sum(v_i(S))))$	45.988
$avg(sg(sum(v_i(S)))))$	40.236
$max(max(v_i(S))))$	15.033
$max(avg(v_i(S))))$	13.687
$max(sum(sv_i(S))))$	13.443
$max(max(sv_i(S))))$	12.029
$max(max(sg(v_i(S)))))$	11.586
$max(avg(sv_i(S))))$	11.559
$max(sum(v_i(S))))$	11.389
$max(sum(sg(v_i(S)))))$	10.451
$max(avg(sg(v_i(S)))))$	7.493

Table 6: Features in Multi Layer Perceptron after merging those that would generate the same heuristic.

5.2. Creating a unified ranking

We now wish to create a combined feature ranking across the four models. We will create this by viewing it as a voting problem, where each of the models has a different ranking for the features and we have to get a combined ranking. A usual method to combine such rankings is the Borda Count: where for each ranking we assign a penalisation of 1 to the first, 2 to the second, and so on; at the end ranking the options from lowest penalization to highest. However, this method strongly penalises being very badly ranked for just one of the models. Meaning that a feature that is the most relevant for three of the models while being the most irrelevant for one of the models would not end up in a good position in the final ranking.

Thus we use a modification, the Dowdall System, that assigns the first feature a reward of 1, the second a reward of 1/2, the third receives 1/3, and so until, at the end ranking the options from higher reward to lowest. See for example the work of Fraenkel and Grofman [24] on these different voting systems. We chose the Dowdall system because it allows merging the preferences without a big penalisation for being a very distant preference for one of the models, contrary to the Borda Count.

Table 7 shows the features and their scores in our final ranking and their rewards from the Dowdall System. The most voted feature is the sum of the degrees of the polynomials in the set, this is precisely the feature used in gmods, which was found in [50] through the study of the complexity analysis of CAD. In this ranking: the third, tenth and eighth most voted features describe the tools used by the Brown heuristic [6].

The second most voted feature, the average across the polynomials of the average degree of the variable across the monomials, is a feature that has been never considered relevant in any prior work.

The top six features are chosen for further experimentation in the next section to find new human-level heuristics. We draw the line here as there is a big drop off and the lower features score more similarly to each other, i.e. the feature scores stabilize after the sixth score. This stabilization of the scores is illustrated in Figure 5.

Feature Name	Voted Score
$sum(max(v_i(S))))$	3.333
$avg(avg(v_i(S))))$	2.167
$sum(sum(v_i(S))))$	1.158
$avg(avg(sg(v_i(S)))))$	1.15
$sum(sg(sum(v_i(S)))))$	0.794
$sum(max(sv_i(S))))$	0.787
$avg(avg(sv_i(S))))$	0.583
$sum(sum(sg(v_i(S)))))$	0.554
$max(sum(v_i(S))))$	0.475
$sum(sum(sv_i(S))))$	0.472
$max(max(v_i(S))))$	0.467
$max(sum(sg(v_i(S)))))$	0.3
$max(avg(v_i(S))))$	0.245
$max(max(sg(v_i(S)))))$	0.218
$max(avg(sg(v_i(S)))))$	0.21
$max(sum(sv_i(S))))$	0.209
$max(avg(sv_i(S))))$	0.192
$max(max(sv_i(S))))$	0.191

Table 7: Voted score of merged and aggregated features across all models.

6. New Human-Level Heuristics Motivated by XAI

In this section, we will first introduce the methodology we use to evaluate a heuristic, and then experiment with new human-level heuristics motivated by Section 5. We will compare these against the current state-of-the-art heuristics for choosing the variable ordering for CAD.

To evaluate the performance of these heuristics two things are needed: a set of meaningful benchmark examples and metrics to quantify the performance of the heuristics on them. We will use the benchmarks and metrics described recently by del Río and England [50] for this purpose.
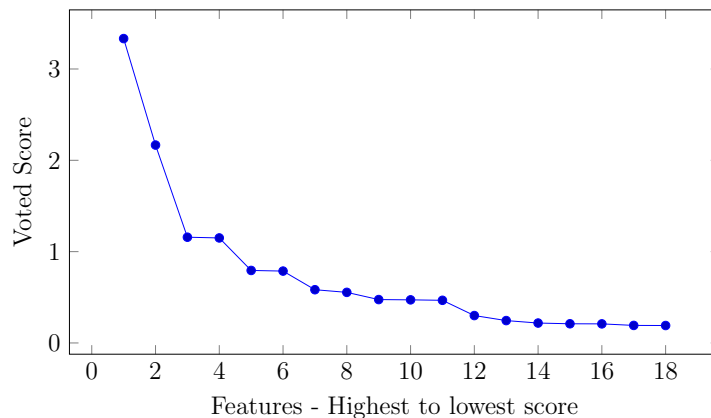
Figure 5: Plot of feature scores

6.1. Benchmarks

We source benchmarks from the three-variable problems in the QF_NRA category of the SMT-LIB [1]. These examples are all satisfiability problems and so do not represent the full application range of CAD which can also address quantifier elimination. However, they do mostly emit from real applications making performance upon them particularly meaningful. Common sources are problems emitting from the theorem prover MetiTarski [46], attempts to prove termination of term-rewrite systems, verification conditions from Keymaera [49], and parametrized generalizations of geometric problems, as well as problems emitting from economics [45] and biology [4].

A sign-invariant CAD is built using Maple [11] for each example in each variable ordering: 5599 problems finished for at least one ordering before the time limit (60 seconds). Of these, the authors select 1019 *"unique"* problems where, as described in [50, Section 4.1], unique means the CAD tree structures built are not identical to another example for every ordering. This benchmark merging is needed since some application sources in the SMT-LIB produce examples that differ from each other only slightly, leading to identical CAD computations, and thus potentially unfair evaluation.

The benchmarks used here are identical to those in [50] and so readers may refer to the open source dataset released for that paper on Zenodo[7].

---

[7]https://doi.org/10.5281/zenodo.6750528

6.2. Evaluation metrics

For evaluation metrics, we again follow del Río and England [50] and use the three selected there as follows.

- `Accuracy:` Percentage of benchmarks in which the optimal ordering (quickest in runtime) coincides with that chosen by the heuristic.

This is the most commonly used tool in ML classification and was used to train the ML models earlier[8]. However, it is not ideally suited to our problem: it considers as equally wrong a terrible ordering and an ordering that is slightly sub-optimal, even if the difference compared to the optimal ordering is so small that it could be due to computational noise.

- `Total time:` Amount of time needed to create a CAD for each benchmark when choosing the variable ordering given by the heuristic.

This metric corresponds most closely to the absolute performance achieved and would be most commonly used to rank a computer algebra implementation. However, it will not take into account the chosen orderings for the smaller problems in our dataset and could allow a small number of large problems to distort findings. Hence del Río and England [50] proposed the next metric to address the shortcomings of the above: a metric that does not ignore the simple problems, but at the same time distinguishes between different non-optimal orderings.

- `Markup:` The markup of the time needed by the heuristic choice with respect to using the best choice. Defined as the average across all benchmarks of
$$\text{markup} = \frac{t_{heuristic} - t_{optimal}}{t_{optimal} + 1}.$$

The $+1$ is added in the denominator to reduce the impact that the computational noise can have on problems that only require tenths of seconds.

When a CAD could not be created within the time limit using the ordering chosen by a heuristic, the time of the heuristic was set in the dataset to twice the time limit. These runtimes are also in the Zendodo data release for [50].

---

[8]See [22] for an attempt to train ML partially with a different metric.

6.3. Existing heuristics

The prior state-of-the-art heuristic for choosing CAD variable ordering is gmods, as presented by del Río and England [50]. This is based on a metric motivated by the complexity analysis of CAD. The analysis found that given a set of polynomials S_n and the ordering $x_{i_n} \succ \cdots \succ x_{i_1}$, the maximum number of cells that can be generated is proportional to

$$\prod_{j=1}^{n} sum(max(v_{i_j}(S_j))), \tag{1}$$

where S_{j-1} is the projection of S_j with respect to x_{i_j}.

The heuristic mods was first proposed by del Río and England [50] as a heuristic that computes the sets $S_1, \ldots, S_{n-1}$ for each variable ordering and then chooses the ordering that minimises the worst case complexity, i.e. that minimises (1). This performed very well, in the sense that it chooses good orderings. However, it is an expensive heuristic and the cost of computing the CAD projection phase for all the orderings was not compensated by the improved quality of the ordering it chooses for the dataset in question. This is illustrated later by the experimental results in Table 8, where mods took into account the additional cost of the heuristic while free-mods does not.

To solve this issue del Río and England [50] decided to avoid doing projections for choosing the variable ordering that would not be utilised by the CAD. Instead, they proposed to pick the first variable of the ordering x_{i_n} to be the one with minimal $sum(max(v_i(S_n)))$, choosing the one with the lowest index in case of a tie. Only then do they take a CAD projection of S_n with respect to x_{i_n} to obtain S_{n-1}, which is then studied to pick the second variable $x_{i_{n-1}}$ as the one with minimal $sum(max(v_i(S_{n-1})))$. This allows the heuristic to gather as much information as possible before choosing the next variable, continuing this process to determine the whole variable ordering.

This heuristic does not require any more CAD projections than those needed to run the CAD algorithm. In this paper we refer to heuristics that follow this structure as "*greedy*" (in the traditional algorithmic sense that they make a choice based on the local information). All the heuristics introduced below are greedy[9].

---

[9]We note that greedy-sotd presented in [17] is not greedy in the sense of this paper because it makes a decision based on all the possible single projections for the next step (still less than the original sotd which computed all full projections).

Since the source benchmarks contain a bias towards certain orderings, the authors choose a variable randomly whenever there is a tie, instead of choosing the one with the lowest index or the first lexicographically which could allow that bias to effect the results.

For many years, a well-known and widely used heuristic is that of Brown [6]. Denoted `Brown`, this uses three criteria, in turn, breaking ties with the subsequent ones. It projects first the variable that minimizes:

- $max(max(v_i(S)))$, (i.e. variable with lowest overall degree); breaking ties with

- $max(max(sv_i(S)))$, (i.e. variable with the lowest total degree in the monomials containing it); breaking ties with

- $sum(sum(v_i(S)))$, (i.e. variable with the lowest sum of all its degrees).

It is not specified what to do if there is a third tie: our implementation simply picks between the tied variables randomly. Moreover, it is not specified whether the whole variable ordering is chosen at once or one variable at a time. In this paper, we implement the greedy variant: i.e. we choose the first variable to be projected, and then use the projected polynomials to choose the next variable in the ordering.

Table 8 evaluates these existing heuristics on our benchmark set. We include also the performance of a virtual best heuristic (the hypothetical heuristic which always picks the optimal ordering) to show the best possible figures for the dataset, and a heuristic that picks at random. As in [50] we find `gmods` to be the best of all metrics except accuracy: where it is beaten by `mods` but at a cost not worth paying for this dataset.

Note that for some heuristics the number of problems completed is not a natural number. This is because the heuristic is not deterministic (die to the random choices to break ties) and the experiments are repeated 5000 times to obtain an average to account for this.

Name	Accuracy	Total time	Markup	# Completed
`gmods`	0.563	**7192.226**	**0.212**	**982.622**
`mods`	**0.639**	8137	0.566	979
`free-mods`	0.639	6637	0.127	990
`Brown`	0.553	7842.564	0.278	968.852
`virtual-best`	1	4822	0	1019
`random`	0.167	20797	4.034	262.5

Table 8: Evaluation of the existing heuristics to choose the variable orderings for CAD. In bold, the best measure of the metric out of all the heuristics.

6.4. Creating new greedy heuristics from XAI identified features

We will now build new heuristics solely based on a single feature from the earlier ML models (a category that `gmods` itself falls in). We aim to see how powerful these features are when taking decisions on their own.

The explainability analysis in Section 5 has found $sum(max(v_i(S)))$ to be the most impactful feature. We will create a greedy heuristic `SumMaxV` based on this feature: it will select the variable for which this feature is minimal for each projection (and in the case of a tie one of the tied variables will be chosen randomly).

Consider for example the set $S_3 = \{x_1x_2x_3 - 1, x_1^2 - x_2^2x_3\}$; then `SumMaxV` will work as follows. First, the feature will be computed for each variable:

$$sum(max(v_1(S_3))) = sum([1, 2]) = 3,$$
$$sum(max(v_2(S_3))) = sum([1, 2]) = 3,$$
$$sum(max(v_3(S_3))) = sum([1, 1]) = 2.$$

So x_3 is projected first because it has the minimal value. The Lazard projection [42] of S_3 with respect to x_3 gives $\{x_1^3x_2 - x_2^2, x_1x_2, x_2^2, x_1^2\}$. Factorizing and removing repeated factors gives $S_2 = \{x_1^3 - x_2, x_1, x_2\}$ and again we compute this feature for each of the remaining variables:

$$sum(max(v_1(S_2))) = sum([3, 1, 0]) = 4 \text{ and,}$$
$$sum(max(v_2(S_2))) = sum([1, 0, 1]) = 2.$$

So x_2 is projected next because it has the minimal value. Thus heuristic `SumMaxV` chooses the ordering $x_3 \succ x_2 \succ x_1$.

6.5. Performance of new single feature heuristics

We now evaluate the new heuristics based upon the features that ranked in the top six in Table 7 (selected as discussed at the end of Section 5.2. `AvgAvgSgV` is based upon feature $avg(avg(sg(v_i(S))))$, `SumMaxSV` is based upon $sum(max(sv_i(S)))$ and so on. The results are shown in Table 9.

Name	Accuracy	Total time	Markup	# Completed
SumMaxV	**0.563**	7192.226	**0.212**	982.622
AvgAvgV	0.544	**7138.664**	0.224	**983.468**
SumSumV	0.549	7524.758	0.261	975.34
AvgAvgSgV	0.535	8682.616	0.559	956.324
SumSgSumV	0.45	10836.654	1.223	922.456
SumMaxSV	0.509	8771.652	0.563	956.45

Table 9: Evaluation metrics for the new heuristics to choose the variable orderings for CAD. In bold, the best measure of the metric out of all the heuristics.

As can be observed in the table there are two features that stand out: `SumMaxV` and `AvgAvgV`. They achieve similar results with the best dependent on the metric used and none of them dominating the other.

6.6. New heuristics of triples of features identified by SHAP

We have now seen a number of single-feature heuristics doing well, without a clear winner. We will next combine these into triples, using subsequent ones as tie-breakers, similarly to the heuristic of Brown [6]. This should give a more informed choice than selecting a random variable and is still certainly at a human level in size and complexity.

We will denote these heuristics based on multiple features by concatenating their names with a >, where smaller indicates that it is used as a tie-breaker for the bigger features. Following this notation the `Brown` heuristic would be denoted as `MaxMaxV>MaxMaxSV>SumSumSgV`.

There are 120 possible ordered triples that can be created from the top six features of Table 7. For each we create a heuristic that chooses the variable that minimises the first feature, breaking ties using the second feature and the third if necessary, and finally breaking ties by choosing a random variable if there is a tie for all three features.

Performance statistics for the two triples that are best in one of the metrics are given in Table 10, along with `Brown` that makes choices with a similar quantity of information.

Name	Accuracy	Total time	Markup	# Completed
Brown	0.553	7842.564	0.278	968.852
T1	0.567	**6896.276**	0.193	**985.743**
T2	**0.583**	6896.672	**0.188**	984.773

Table 10: Evaluation metrics for the different heuristics to choose the variable orderings for CAD. In bold, the best measure of the metric out of all the heuristics. T1=SumMaxV>AvgAvgV>SumSumV and T2=SumMaxV>AvgAvgSgV>SumSumV

Table 10 shows that the heuristic T1=SumMaxV>AvgAvgV>SumSumV wins in two of the four metrics: total time and the number of problems completed within the time limit. Note that this heuristic has been created from the top three features XAI suggested features, in the order they appear in the voting system.

- $sum(max(v_i(S)))$, (i.e. variable with the lowest degree in the product of the polynomials);
  breaking ties with

- $avg(avg(v_i(S)))$, (i.e. variable with the lowest average of average degree in the polynomials); breaking ties with

- $sum(sum(v_i(S)))$, (i.e. variable with the lowest sum of all its degrees).

This triplet also outperforms all the singleton heuristics.

In Figure 6 we give a survival plot comparing the triplet T1 to the prior-state of the art and the virtual best for just the hardest problems in the dataset (those taking more than 30 seconds in all orderings). Observe the leap that exists between the prior state-of-the-art and the best possible in theory, with the new XAI informed triplet bridging a good deal of that gap.

7. Conclusions

In this paper, we have devised a new state-of-the-art human-level heuristic for choosing a CAD variable ordering. Previous approaches to the problem were of two types: either an expert handpicking features to create a human-designed heuristic; or an ML model computing with many features under the hood. We presented an intermediate approach where we start with the ML model but then use XAI to select important features for later use independent on any ML technology.
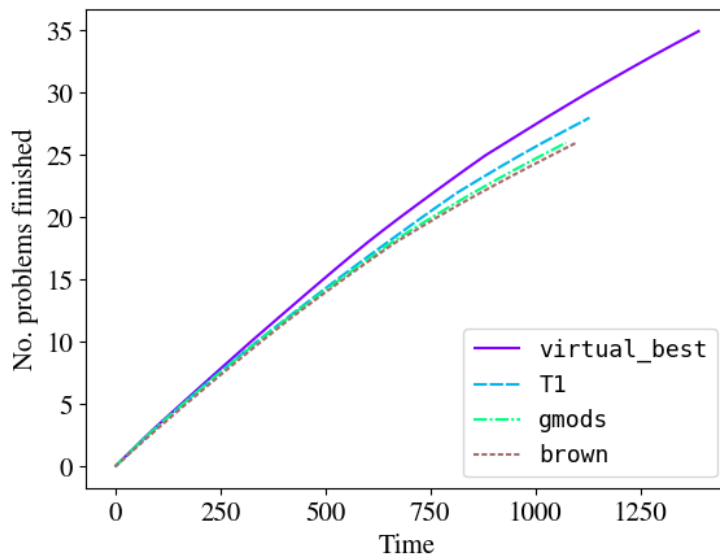
Figure 6: Survival plot for best-existing heuristics and new heuristics proposed in the problems that take more than 30 seconds for all the orderings.

We note that SHAP was able to find the feature identified recently by del Río and England [50] that was the prior state-of-the-art heuristic; and identified also the tools of the heuristic in Brown [6]. But it also found an important feature not previously considered by experts.

There is of course scope for future work. Using multiple features clearly improved the results, and there are certainly others ways we could consider combining them. Perhaps this combination itself could also be ML informed? It would also be interesting to explore ML models that are inherently explainable, rather than relying only on the post-hoc explainability analysis.

The experiments presented in this paper are for selecting CAD variable ordering. However, the methodology could be applied almost directly to variable ordering choices for other algorithms, and could be adapted to many algorithm optimisation problems in symbolic computation and beyond. We have shown that ML has a role to play in symbolic computation, even for developers who prefer not to include a reliance upon it in their final code. It can be used as a tool to guide and develop algorithms at a human level. We thus hope and expect to see other applications of XAI to inform decisions in CASs and optimise symbolic computation algorithms in the coming years.

Acknowledgements

References

[1] Barrett, C., Fontaine, P., Tinelli, C., . The Satisfiability Modulo Theories Library (SMT-LIB). Online Resource. URL: https://smtlib.cs.uiowa.edu/.

[2] Bernal, E.A., Hauenstein, J.D., Mehta, D., Regan, M.H., Tang, T., 2023. Machine learning the real discriminant locus. Journal of Symbolic Computation 115, 409–426. URL: https://doi.org/10.1016/j.jsc.2022.08.001.

[3] Bishop, C.M., 2006. Pattern Recognition and Machine Learning. Springer. URL: https://link.springer.com/book/9780387310732.

[4] Bradford, R., Davenport, J.H., England, M., Errami, H., Gerdt, V., Grigoriev, D., Hoyt, C., Košta, M., Radulescu, O., Sturm, T., Weber, A., 2020. Identifying the parametric occurrence of multiple steady states for some biological networks. Journal of Symbolic Computation 98, 84–119. URL: https://doi.org/10.1016/j.jsc.2019.07.008.

[5] Bradford, R., Davenport, J.H., England, M., Wilson, D., 2013. Optimising problem formulations for cylindrical algebraic decomposition, in: Carette, J., Aspinall, D., Lange, C., Sojka, P., Windsteiger, W. (Eds.), Intelligent Computer Mathematics. Springer Berlin Heidelberg. volume 7961 of *Lecture Notes in Computer Science*, pp. 19–34. URL: http://dx.doi.org/10.1007/978-3-642-39320-4_2.

[6] Brown, C.W., 2004. Companion to the tutorial "Cylindrical algebraic decomposition", presented at ISSAC '04. URL: https://www.usna.edu/Users/cs/wcbrown/research/ISSAC04/handout.pdf.

[7] Brown, C.W., 2015. Open non-uniform cylindrical algebraic decompositions, in: Proceedings of the 2015 International Symposium on Symbolic and Algebraic Computation, ACM. pp. 85–92. URL: https://doi.org/10.1145/2755996.2756654.

[8] Brown, C.W., Davenport, J.H., 2007. The complexity of quantifier elimination and cylindrical algebraic decomposition, in: Proceedings of the 2007 International Symposium on Symbolic and Algebraic Computation, ACM. pp. 54–60. URL: https://doi.org/10.1145/1277548.1277557.

[9] Buchberger, B., 2006. Bruno Buchberger's PhD thesis (1965): An algorithm for finding the basis elements of the residue class ring of a zero dimensional polynomial ideal. Journal of Symbolic Computation 41, 475–511. URL: https://doi.org/10.1016/j.jsc.2005.09.007.

[10] Carette, J., 2004. Understanding expression simplification, in: Proceedings of the 2004 International Symposium on Symbolic and Algebraic Computation, ACM. pp. 72–79. URL: https://doi.org/10.1145/1005285.1005298.

[11] Chen, C., Moreno Maza, M., Xia, B., Yang, L., 2009. Computing cylindrical algebraic decomposition via triangular decomposition, in: Proceedings of the 2009 International Symposium on Symbolic and Algebraic Computation, ACM. pp. 95–102. URL: https://doi.org/10.1145/1576702.1576718.

[12] Chen, C., Zhu, Z., Chi, H., 2020. Variable ordering selection for cylindrical algebraic decomposition with artificial neural networks, in: Bigatti, A., Carette, J., Davenport, J.H., Joswig, M., de Wolff, T. (Eds.), Mathematical Software – ICMS 2020, Springer International Publishing. pp. 281–291. URL: https://doi.org/10.1007/978-3-030-52200-1_28.

[13] Collins, G.E., 1975. Quantifier elimination for real closed fields by cylindrical algebraic decomposition. Lecture Notes in Computer Science 33, 134–183. URL: https://doi.org/10.1007/3-540-07407-4_17.

[14] Collins, G.E., Hong, H., 1991. Partial cylindrical algebraic decomposition for quantifier elimination. Journal of Symbolic Computation 12, 299–328. URL: https://doi.org/10.1016/S0747-7171(08)80152-6.

[15] Davenport, J.H., Heintz, J., 1988. Real quantifier elimination is doubly exponential. Journal of Symbolic Computation 5, 29–35. URL: https://doi.org/10.1016/S0747-7171(88)80004-X.

[16] Davies, A., Veličković, P., Buesing, L., Blackwell, S., Zheng, D., Tomašev, N., Tanburn, R., Battaglia, P., Blundell, C., Juhász, A., Lackenby, M., Williamson, G., Hassabis, D., Kohli, P., 2021. Advancing mathematics by guiding human intuition with AI. Nature 600, 70–74. URL: https://doi.org/10.1038/s41586-021-04086-x.

[17] Dolzmann, A., Seidl, A., Sturm, T., 2004. Efficient projection orders for CAD, in: Proceedings of the 2004 International Symposium on Symbolic and Algebraic Computation, ACM. pp. 111–118. URL: https://doi.org/10.1145/1005285.1005303.

[18] Du, M., Liu, N., Hu, X., 2019. Techniques for interpretable machine learning. Commun. ACM 63, 68–77. URL: https://doi.org/10.1145/3359786.

[19] England, M., Bradford, R., Davenport, J.H., 2020. Cylindrical algebraic decomposition with equational constraints. Journal of Symbolic Computation 100, 38–71. URL: https://doi.org/10.1016/j.jsc.2019.07.019.

[20] England, M., Florescu, D., 2019. Comparing machine learning models to choose the variable ordering for cylindrical algebraic decomposition, in: Kaliszyk, C., Brady, E., Kohlhase, A., Sacerdoti, C.C. (Eds.), Intelligent Computer Mathematics, Springer International Publishing. pp. 93–108. URL: https://doi.org/10.1007/978-3-030-23250-4_7.

[21] Florescu, D., England, M., 2019. Algorithmically generating new algebraic features of polynomial systems for machine learning, in: Abbott, J., Griggio, A. (Eds.), Proceedings of the 4th Workshop on Satisfiability Checking and Symbolic Computation (SC^2 2019), p. 12. URL: http://ceur-ws.org/Vol-2460/.

[22] Florescu, D., England, M., 2020a. Improved cross-validation for classifiers that make algorithmic choices to minimise runtime without

compromising output correctness, in: Slamanig, D., Tsigaridas, E., Zafeirakopoulos, Z. (Eds.), Mathematical Aspects of Computer and Information Sciences (Proc. MACIS '19), Springer International Publishing. pp. 341–356. URL: https://doi.org/10.1007/978-3-030-43120-4_27.

[23] Florescu, D., England, M., 2020b. A machine learning based software pipeline to pick the variable ordering for algorithms with polynomial inputs, in: Bigatti, A., Carette, J., Davenport, J.H., Joswig, M., de Wolff, T. (Eds.), Mathematical Software – ICMS 2020, Springer International Publishing. pp. 302–322. URL: https://doi.org/10.1007/978-3-030-52200-1_30.

[24] Fraenkel, J., Grofman, B., 2014. The Borda Count and its real-world alternatives: Comparing scoring rules in Nauru and Slovenia. Australian Journal of Political Science 49, 186–205. URL: https://doi.org/10.1080/10361146.2014.900530, doi:10.1080/10361146.2014.900530.

[25] Giovini, A., Mora, T., Niesi, G., Robbiano, L., Traverso, C., 1991. One sugar cube, please; or selection strategies in the Buchberger algorithm, in: Proceedings of the 1991 International Symposium on Symbolic and Algebraic Computation, ACM. pp. 49–54. URL: https://doi.org/10.1145/120694.120701.

[26] Gunning, D., Aha, D., 2019. DARPA's Explainable Artificial Intelligence (XAI) Program. AI Magazine 40, 44–58. URL: https://doi.org/10.1609/aimag.v40i2.2850. number: 2.

[27] Hauenstein, J.D., He, Y.H., Kotsireas, I., Mehta, D., Tang, T., 2023. Special issue on algebraic geometry and machine learning. Journal of Symbolic Computation 118, 93–94. URL: https://doi.org/10.1016/j.jsc.2022.10.003.

[28] He, Y.H., 2022. Machine-learning mathematical structures. International Journal of Data Science in the Mathematical Sciences 1, 1–25. URL: https://doi.org/10.1142/S2810939222500010.

[29] Herm, L.V., Heinrich, K., Wanner, J., Janiesch, C., 2022. Stop ordering machine learning algorithms by their explainability! A user-centered investigation of performance and explainability. International Journal

of Information Management 69, 102538. URL: https://doi.org/10.1016/j.ijinfomgt.2022.102538.

[30] Huang, Z., England, M., Davenport, J.H., Paulson, L., 2016. Using machine learning to decide when to precondition cylindrical algebraic decomposition with Groebner bases, in: 18th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC '16), IEEE. pp. 45–52. URL: https://doi.org/10.1109/SYNASC.2016.020.

[31] Huang, Z., England, M., Wilson, D., Davenport, J.H., Paulson, L., Bridge, J., 2014. Applying machine learning to the problem of choosing a heuristic to select the variable ordering for cylindrical algebraic decomposition, in: Watt, S.M., Davenport, J.H., Sexton, A.P., Sojka, P., Urban, J. (Eds.), Intelligent Computer Mathematics. Springer International. volume 8543 of *Lecture Notes in Artificial Intelligence*, pp. 92–107. URL: http://dx.doi.org/10.1007/978-3-319-08434-3_8.

[32] Huang, Z., England, M., Wilson, D.J., Bridge, J., Davenport, J.H., Paulson, L.C., 2019. Using machine learning to improve cylindrical algebraic decomposition. Mathematics in Computer Science 13, 461–488. doi:10.1007/s11786-019-00394-8.

[33] Jovanovic, D., de Moura, L., 2012. Solving non-linear arithmetic, in: Gramlich, B., Miller, D., Sattler, U. (Eds.), Automated Reasoning: 6th International Joint Conference (IJCAR). Springer. volume 7364 of *Lecture Notes in Computer Science*, pp. 339–354. URL: https://doi.org/10.1007/978-3-642-31365-3_27.

[34] Kobayashi, M., Iwane, H., Matsuzaki, T., Anai, H., 2016. Efficient subformula orders for real quantifier elimination of non-prenex formulas, in: Kotsireas, S.I., Rump, M.S., Yap, K.C. (Eds.), Mathematical Aspects of Computer and Information Sciences (MACIS '15), Springer International Publishing. pp. 236–251. URL: https://doi.org/10.1007/978-3-319-32859-1_21.

[35] Kremer, G., Ábrahám, E., 2020. Fully incremental CAD. Journal of Symbolic Computation 100, 11–37. URL: https://doi.org/10.1016/j.jsc.2019.07.018.

[36] Kuipers, J., Ueda, T., Vermaseren, J.A.M., 2015. Code optimization in FORM. Computer Physics Communications 189, 1–19. URL: https://doi.org/10.1016/j.cpc.2014.08.008.

[37] Lample, G., Charton, D., 2020. Deep learning for symbolic mathematics, in: Mohamed, S., White, M., Cho, K., Song, D. (Eds.), Eighth International Conference on Learning Representations (ICLR 2020), p. 24. URL: https://iclr.cc/virtual_2020/poster_S1eZYeHFDS.html.

[38] Lundberg, S., . Waterfall plot — SHAP documentation. URL: https://shap.readthedocs.io/en/latest/example_notebooks/api_examples/plots/waterfall.html.

[39] Lundberg, S., 2022. SHAP v1.40.0 (GitHub). URL: https://github.com/slundberg/shap.

[40] Lundberg, S.M., Lee, S.I., 2017. A unified approach to interpreting model predictions, in: Proceedings of the 31st International Conference on Neural Information Processing Systems, Curran Associates Inc., Red Hook, NY, USA. p. 4768–4777. URL: https://dl.acm.org/doi/10.5555/3295222.3295230.

[41] Manubens, M., Moroz, G., Chablat, D., Rouillier, F., Wenger, P., 2012. Cusp points in the parameter space of degenerate 3-RPR planar parallel manipulators. J. Mechanisms Robotics 4, 041003. URL: https://doi.org/10.1115/1.4006921.

[42] McCallum, S., Hong, H., 2016. On using lazard's projection in cad construction. Journal of Symbolic Computation 72, 65–81. doi:10.1016/j.jsc.2015.02.001.

[43] McCallum, S., Parusińiski, A., Paunescu, L., 2019. Validity proof of Lazard's method for CAD construction. Journal of Symbolic Computation 92, 52–69. URL: https://doi.org/10.1016/j.jsc.2017.12.002.

[44] Molnar, C., 2023. Interpretable Machine Learning. URL: https://christophm.github.io/interpretable-ml-book/.

[45] Mulligan, C.B., Davenport, J.H., England, M., 2018. TheoryGuru: A Mathematica package to apply quantifier elimination technology to economics, in: Davenport, J.H., Kauers, M., Labahn, G., Urban,

J. (Eds.), Mathematical Software – Proc. ICMS 2018, Springer International Publishing. pp. 369–378. URL: https://doi.org/10.1007/978-3-319-96418-8_44.

[46] Paulson, L.C., 2012. Metitarski: Past and future, in: Beringer, L., Felty, A. (Eds.), Interactive Theorem Proving. Springer. volume 7406 of *Lecture Notes in Computer Science*, pp. 1–10. URL: https://doi.org/10.1007/978-3-642-32347-8_1.

[47] 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., 2011. Scikit-learn: Machine learning in Python. Journal of Machine Learning Research 12, 2825–2830. URL: http://www.jmlr.org/papers/v12/pedregosa11a.html.

[48] Peifer, D., Stillman, M., Halpern-Leistner, D., 2020. Learning selection strategies in Buchberger's algorithm, in: Daumé III, H., Singh, A. (Eds.), Proceedings of the 37th International Conference on Machine Learning (ICML 2020), PMLR. pp. 7575–7585. URL: https://proceedings.mlr.press/v119/peifer20a.html.

[49] Platzer, A., Quesel, J., Rümmer, P., 2009. Real world verification, in: Schmidt, R.A. (Ed.), Automated Deduction (CADE-22). Springer Berlin Heidelberg. volume 5663 of *Lecture Notes in Computer Science*, pp. 485–501. URL: https://doi.org/10.1007/978-3-642-02959-2_35.

[50] del Río, T., England, M., 2022. New heuristic to choose a cylindrical algebraic decomposition variable ordering motivated by complexity analysis, in: Boulier, F., England, M., Sadykov, T.M., Vorozhtsov, E.V. (Eds.), Computer Algebra in Scientific Computing. Springer International Publishing. volume 13366 of *Lecture Notes in Computer Science*, pp. 300–317. URL: https://doi.org/10.1007/978-3-031-14788-3_17.

[51] Röst, G., Sadeghimanesh, A., 2021. Exotic bifurcations in three connected populations with allee effect. International Journal of Bifurcation and Chaos 31. URL: https://doi.org/10.1142/S0218127421502023.

[52] Shapley, L.S., 1953. 17. a value for n-person games, in: Contributions to the Theory of Games (AM-28), Volume II, pp. 307–317. URL: https://doi.org/10.1515/9781400881970-018.

[53] Simpson, M.C., Yi, Q., Kalita, J., 2016. Automatic algorithm selection in computational software using machine learning, in: 15th IEEE International Conference on Machine Learning and Applications (ICMLA 2016), pp. 355–360. URL: https://doi.org/10.1109/ICMLA.2016.0064.

[54] Strzeboński, A., 2006. Cylindrical algebraic decomposition using validated numerics. Journal of Symbolic Computation 41, 1021–1038. URL: https://doi.org/10.1016/j.jsc.2006.06.004.

[55] Wilson, D., England, M., Bradford, R., Davenport, J.H., 2015. Using the distribution of cells by dimension in a cylindrical algebraic decomposition. Proceedings - 16th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing, SYNASC 2014 , 53–60doi:10.1109/SYNASC.2014.15.