

---

# WARFARIN DOSE ESTIMATION ON MULTIPLE DATASETS WITH AUTOMATED HYPERPARAMETER OPTIMISATION AND A NOVEL SOFTWARE FRAMEWORK

---

A PREPRINT

**Gianluca Truda**

Department of Computer Science  
University of Cape Town  
Cape Town, South Africa  
trdgia001@myuct.ac.za

**Patrick Marais**

Department of Computer Science  
University of Cape Town  
Cape Town, South Africa  
patrick@cs.uct.ac.za

October 28, 2020

## ABSTRACT

Warfarin is an effective preventative treatment for arterial and venous thromboembolism, but requires individualised dosing due to its narrow therapeutic range and high individual variation. Many machine learning techniques have been demonstrated in this domain. This study evaluated the accuracy of the most promising algorithms on the International Warfarin Pharmacogenetics Consortium dataset and a novel clinical dataset of South African patients. Support vectors and linear regression were amongst the top performers in both datasets and performed comparably to recent stacked ensemble approaches, whilst neural networks were one of the worst performers in both datasets. We also introduced genetic programming to automatically optimise model architectures and hyperparameters without human guidance. Remarkably, the generated models were found to match the performance of the best models hand-crafted by human experts. Finally, we present a novel software framework (Warfit-learn) for warfarin dosing research. It leverages the most successful techniques in preprocessing, imputation, and parallel evaluation, with the goal of accelerating research and making results in this domain more reproducible.

**Keywords** Warfarin · machine learning · genetic programming · Python · supervised learning · anticoagulant · pharmacogenetics · software

## 1 Introduction

Many individuals suffer from blood clots that lead to arterial and venous thromboembolism. The standard method for treating these conditions is the use of anticoagulant drugs such as vitamin K antagonists, the most widely used of which is warfarin. Whilst effective, the drug has a narrow therapeutic range and severe side-effects at extreme concentrations. This makes the precise dosing of warfarin an important concern for clinicians. Unfortunately, warfarin metabolism differs across individuals based on age, weight, genetics, diet, drug interactions, and various pre-existing conditions [1, 2]. To standardise the process of anticoagulant monitoring, the World Health Organisation adopted the international

normalised ratio (INR) [3]. INR has become the standard measurement for anticoagulation monitoring around the world [4].

Many studies [5, 6, 7, 8, 9, 10, 11, 12] have looked at applying statistical models to the problem of individualised warfarin dosing. Accurate models improve the ability of clinicians to prescribe the correct warfarin doses to their patients, whilst minimising the time required to do so. They also reduce the risk of severe haemorrhaging in patients and the number of visits required to establish a therapeutic dose [13, 14, 15]. Unfortunately, warfarin datasets are small and noisy, requiring the use of specialised data transformations and highly-optimised learning algorithms. Breakthrough techniques are of significance to the medical research community and could lead to improvements in decision support for those administering warfarin therapy.

The learning algorithm chosen to produce the dosing model has a notable impact on accuracy and robustness. Moreover, the precise hyperparameters chosen for the algorithm can drastically affect performance.

Unfortunately, replicating the precise methodology of previous studies has proven difficult, as the methods for preprocessing, imputation, data stratification, and evaluation vary. Moreover, these methods are typically described textually, with open-sourced software implementations being unavailable prior to this study. We introduce a new software framework (Warfit-learn) that implements the best available techniques. We validate the framework by using it to compare a variety of established and novel algorithms on two different warfarin datasets.

## 1.1 Related work

### 1.1.1 International warfarin pharmacogenetics consortium (IWPC)

The seminal work in warfarin dose estimation with pharmacogenetic data made use of 5052 patient records from the initial version of the dataset utilised by later studies [16]. Researchers split the dataset randomly with an 80/20 ratio into derivation and validation sets. The study selected the two key metrics that would come to be used in almost all future work – mean absolute error (MAE) and percentage of patients within 20% of the therapeutic dose (PW20). They found that a multiple linear regression technique yielded better results than the existing clinical algorithms, with an MAE of  $8.5 \pm 1.7$  mg/week in the validation cohort [16].

### 1.1.2 Liu et al. 2015

A notable study is that of Liu et al. in 2015 [5], which compared the average performance of 9 learning algorithms on the IWPC dataset [17]. They removed patients missing height, weight, age, or genotype data, and patients not at a stable warfarin dose – leaving 4798 patients remaining. Using libraries in the R language, they implemented 9 algorithms. They obtained the average performance of each algorithm in terms of PW20 and MAE with 100 rounds of 80/20 re-sampling from the filtered dataset. The focus of their study was evaluating a range of off-the-shelf algorithms across dosage ranges and racial groups, but their top results (PW20 = 46.35%, MAE = 8.84) in the combined cohort set an initial benchmark for performance on the IWPC data. This result was achieved with a multivariate adaptive regression splines algorithm, though the authors note that Bayesian additive regression trees and support vector machines were also high-performing algorithms.

### 1.1.3 Ma et al. 2018

Another notable study is that of Ma et al. in late 2018 [12], which made advances using ensemble techniques and improved imputation of missing data. Utilising the same IWPC dataset [17], Ma et al. imputed missing weight and height values using a linear regression model. For imputing height, the variables were weight, race, and sex. For imputing weight, they were height, race, and sex. They imputed missing values for the *VKORC1 rs9923231* genotype using the IWPC’s formula [18] based on race and linkage disequilibrium in *VKORC1*. Even after excluding outliers, 5743 subjects remained in the data set, resulting in a cohort nearly 20% greater than that of Liu et al. The raw warfarin

dose was square rooted to compensate for the skewed distribution of the variable. The researchers made use of a stacked generalisation (stacking) framework to implement heterogeneous ensembles. Results were obtained through 100 rounds of 80/20 re-sampling. Their findings suggest that stacked generalisation ensembles are significantly more accurate than the existing isolated models on the IWPC dataset. Crucially, however, these high-performing stacked ensembles were given extra input parameters that were not accessible to the other algorithms. In addition to the 11 common parameters used in the IWPC and Liu et al. models [17, 5], their stacked ensembles included indicators for diabetes mellitus, heart failure, valve replacement, use of statins, and additional *VKORC1* genotypes. Despite acknowledging the use of additional parameters in their methodology, the authors still claim to have performed a direct comparison. Although they appear to have successfully utilised these additional parameters, the inconsistency of the methodology renders their findings inconclusive. Considering that hyperparameter tuning was performed on the entire dataset, it is hard to differentiate their reported accuracy improvements from possible overfitting effects. With our open-sourced Warfit-learn framework, we hope to reduce the possibility of such inconsistencies in future work.

## 1.2 Model optimisation

The most common approach to model development is for an expert to repeatedly train models with various learning algorithms and tweak their hyperparameters until performance is maximised. Ensemble methods – such as bagging, boosting, and stacking – combine the outputs from several models to obtain better overall performance on unseen data. This has been found to increase model performance in a number of warfarin dosing studies [5, 19, 8, 12]. Another approach is the use of automatic machine learning (autoML), which employs meta-algorithms to automate the task of optimisation.

In the past, autoML approaches have focused on optimising subsets of the machine learning pipeline [20]. Grid search is a common form of hyperparameter optimisation based on a brute-force search of algorithm hyperparameters, with the drawback of computational complexity. Conversely, random evaluation within the search space typically discovers an effective hyperparameter set more quickly than exhaustive search [21]. A trade-off between the expensive (but comprehensive) exhaustive search and the efficient (but uncertain) random evaluation is desired. A well-known approach to this trade-off is evolutionary algorithms. Indeed, the use of genetic programming – a subset of evolutionary algorithms – has been found to produce better models than human experts in a number of tasks [22, 23, 24, 25, 26, 27]. This promising avenue to autoML was explored in our paper.

## 1.3 Genetic programming

Genetic programming emulates the process of natural selection as an optimisation strategy. In the context of this study, the smallest components are machine learning algorithms and data processors. Each is encoded as a gene, with parameters changing according to defined mutation and crossover rates. As with biological evolution, successful genes propagate through the population and the most successful combinations of those genes seed the next generation [28]. This increases performance over time. By simulating many generations with well-chosen evolutionary hyperparameters and population sizes, novel meta-algorithms with high accuracy can be produced. Genetic programming has been shown to develop effective systems in a number of mathematical and computational domains [24, 25, 26, 27].

This study made use of the open-sourced Tree-based Pipeline Optimisation Tool (TPOT) to generate high-performing models through genetic programming [23]. The TPOT framework was developed atop the Distributed Evolutionary Algorithms in Python framework [29]. Each individual in the population is a tree-based machine learning pipeline. Each node in the tree is a pipeline operator, which can be (1) a preprocessor, (2) a decomposer, (3) a feature selector, or (4) a supervised learning algorithm. All of these operators are based on Scikit-learn [30] implementations. These pipelines handle the data from feature extraction through to hyperparameter optimisation. The pipeline operators (and their parameters) are encoded as a gene sequence, which evolves through the genetic programming process [23].

To prevent a tendency toward overly-complex pipelines, TPOT makes use of Pareto optimisation to optimise two objectives simultaneously – maximising pipeline fitness, whilst minimising pipeline complexity (measured as the number of operators). Over many generations of evolution, TPOT’s implementation of the genetic programming algorithm evolves the pipelines by adding new pipeline operators that improve fitness, and by removing redundant or detrimental pipeline operators. The fittest overall pipeline is saved as the representative pipeline at the end of a TPOT run [23]. In other domains, TPOT has been shown to produce a significant improvement over basic machine learning methods, with little involvement from users [23]. This study utilised TPOT to automatically evolve optimal architectures and hyperparameters for warfarin dose estimation. The resulting estimators performed on par with those manually optimised by machine learning experts.

## 2 Materials and methods

An illustration of the study’s methodology can be found in Figure 1 and highlights where our Warfit-learn software framework was applied.

### 2.1 Datasets

Two datasets of warfarin records were used for this study. The standard IWPC dataset was used primarily for comparing new techniques to those in existing literature, whilst a novel dataset from South Africa was used to evaluate how well estimators generalised to a different clinical context and parameter set. Both datasets have similar distributions of weekly warfarin dose, INR, and age (see Table 1).

The International Warfarin Pharmacogenetics Consortium (IWPC) dataset [17] of 6256 patients has been used in a number of notable studies [5, 7, 16, 12] and is the standard reference point for new approaches to automated warfarin dosing. The dataset was compiled collaboratively and includes data from 22 research groups from 9 countries [16]. The dataset is publicly available at [pharmgkb.org/downloads](http://pharmgkb.org/downloads).

The novel dataset was provided by PathCare, a private pathology group in South Africa. Unlike the IWPC dataset, no pharmacogenetic data is available and very limited clinical data is present. A total of 4621 patients from the dataset achieved their target INR range and were considered in this study (see Table 1). As with the IWPC dataset, patients were de-identified prior to distribution. To guarantee the confidentiality of the data, only the authors had access to it during the study. Ethical approval was obtained from the Science Research Ethics Committee at the University of Cape Town and the project was sanctioned by the PathCare Research Committee. To protect the patients, the data is not publicly available.

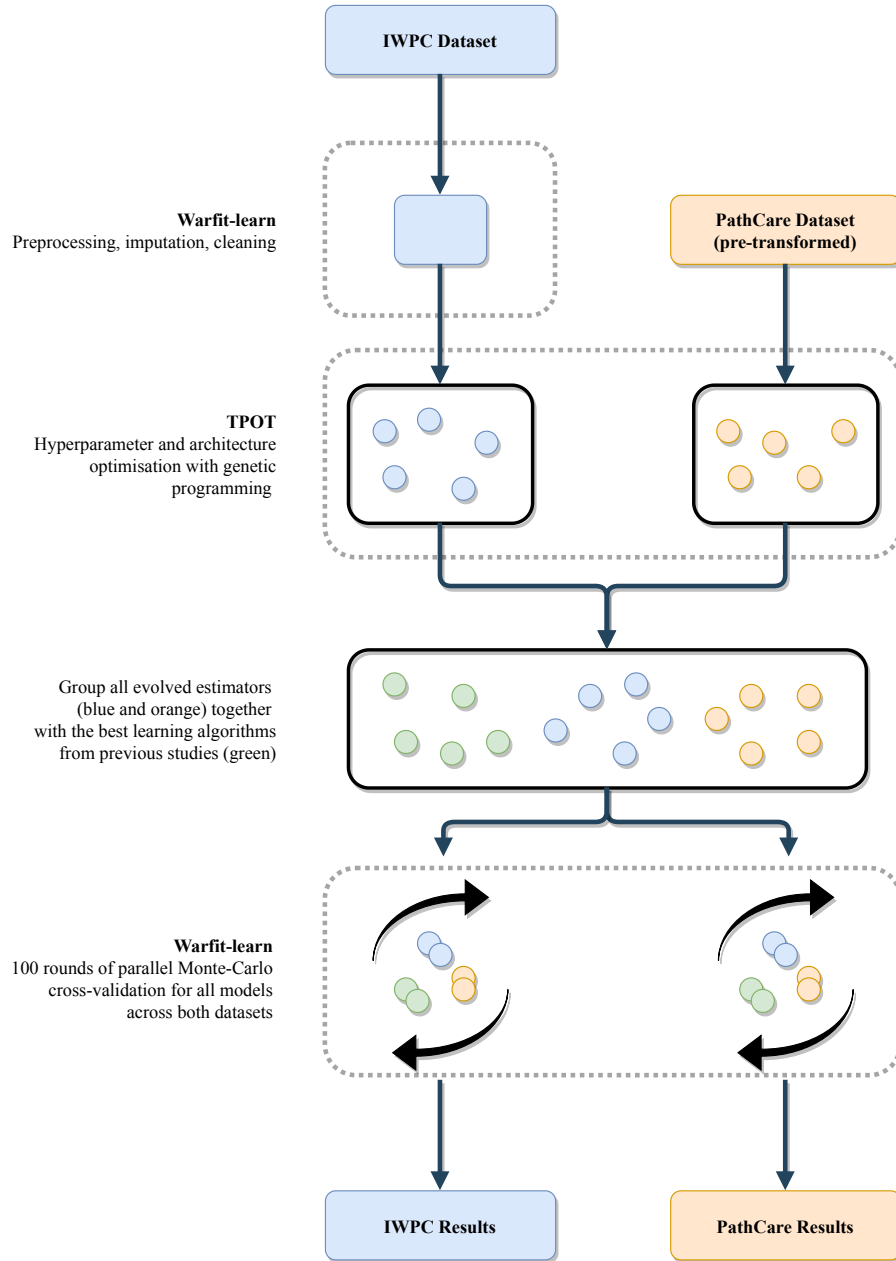


Figure 1: An illustration of the methodology applied in this study, highlighting where our Warfit-learn software framework was utilised and how a collection of estimators were evaluated across multiple datasets.

Table 1: Analysis of the characteristics of the International Warfarin Pharmacogenetics Consortium (IWPC) subjects and subjects from a novel South African dataset (PathCare) included in this study.

	<b>IWPC</b>		<b>PathCare</b>	
Size	5741		4621	
<b>Warfarin dose (mg/week)</b>				
Mean	31.97		29.63	
Median	28.0		27.50	
Range	2.1 – 230.0		0.0 – 210.0	
Standard deviation	16.77		16.04	
<b>Therapeutic INR</b>				
Mean	2.15		2.46	
Median	2.34		2.40	
Range	0.0 – 6.1		1.8 – 3.9	
Standard deviation	0.82		0.33	
<b>Age</b>	<b>Count</b>	<b>%</b>	<b>Count</b>	<b>%</b>
10 - 19	15	0.3	2	0.0
20 - 29	130	2.3	31	0.7
30 - 39	249	4.3	106	2.3
40 - 49	580	10.1	276	6.0
50 - 59	1133	19.7	490	10.6
60 - 69	1401	24.4	942	20.4
70 - 79	1539	26.8	1143	24.7
80 - 89	660	11.5	1127	24.4
90+	34	0.6	204	4.4
<b>Sex</b>	<b>Count</b>	<b>%</b>	<b>Count</b>	<b>%</b>
Females	2466	43.0	2443	52.9
Males	3275	57.0	2178	47.1
Missing	0	0	0	0

Note: These are the cohorts after the preprocessing and imputation stages.

## 2.2 Data cleaning

For the IWPC dataset, this study replicated the work of Ma et al. [12] for filtering the patient records into a viable cohort. This was done to verify their approach and to provide a consistent methodology for performance comparisons. The process is detailed in their paper and in the source code of this study. In this domain, missing data has typically been handled by either dropping whole records [5] or imputing missing values [12, 11, 16]. Because the IWPC data is limited in size, dropping missing values leads to significant data loss and is detrimental to performance. This study leveraged the most successful imputation techniques from previous work – missing values for *VKORC1 rs9923231* were imputed using the IWPC’s formula [18] based on race and linkage disequilibrium, whilst missing height and weight were imputed with a regression model as per the work of Ma et al. [12]. This process is built into the Warfit-learn preprocessing module.

## 2.3 Parameter selection

For the purposes of direct comparison to baseline, this study used the same parameter set defined by the IWPC [17] and utilised by Liu et al. [5]. For the PathCare dataset, three parameter sets were evaluated using standard regression techniques. The top-performing parameter set was selected and used for further evaluation. This set included the sex and age of the patients; as well as records of their aspirin, paracetamol, and amiodarone use; and their history of atrial fibrillation, deep vein thrombosis, and heart-valve replacement. All categorical parameters were vectorised into sparse format so that the resulting feature set was purely numerical.

Table 2: Continued: Analysis of the characteristics of the International Warfarin Pharmacogenetics Consortium (IWPC) subjects and subjects from a novel South African dataset (PathCare) included in this study.

	IWPC		PathCare	
<b>Height (cm)</b>				
Mean	167.98		-	
Median	167.64		-	
Range	125 – 202		-	
Standard deviation	10.58		-	
<b>Weight (kg)</b>				
Mean	78.79		-	
Median	76.0		-	
Range	30 – 238		-	
Standard deviation	22.32		-	
<b>Race (OMB)</b>				
	<b>Count</b>	<b>%</b>	<b>Count</b>	<b>%</b>
White	3095	53.9	-	-
Asian	1515	26.4	-	-
Black	665	11.6	-	-
Missing / mixed race	466	8.1	-	-
<b>VKORC1 Genotype (Imputed)</b>				
	<b>Count</b>	<b>%</b>	<b>Count</b>	<b>%</b>
A/A	1695	29.5	-	-
A/G	2058	35.8	-	-
G/G	1884	32.8	-	-
Unknown	104	1.8	-	-
<b>CYP2C9 Genotype</b>				
	<b>Count</b>	<b>%</b>	<b>Count</b>	<b>%</b>
*1/*1	4230	73.7	-	-
1/*2	755	13.2	-	-
1/*3	482	8.4	-	-
1/*5	0	0.0	-	-
1/*6	0	0.0	-	-
2/*2	58	1.0	-	-
2/*3	68	1.2	-	-
3/*3	20	0.3	-	-
Unknown	128	2.2	-	-
<b>Indications</b>				
	<b>Count</b>	<b>%</b>	<b>Count</b>	<b>%</b>
Enzyme inducer*	61	1.1	-	-
Amiodarone	280	4.9	131	2.8
Smoker	482	8.4	-	-

Notes: These are the cohorts after the preprocessing and imputation stages. The OMB racial categories were used in the IWPC dataset and are defined by the 1997 Office of Management and Budget (OMB) standards for the U.S. Census Bureau. \* = carbamazepine, phenytoin, rifampin, rifampicin.

## 2.4 Preprocessing

The input features were scaled prior to training using the Scikit-learn implementation of `StandardScaler` – which scales the data to unit variance. This desensitised the algorithms to the magnitudes of features. In both datasets, the weekly dose (in mg) was transformed to its square root in order to correct for the skewed distributions. This was then set as the target feature during training. During evaluation, the model predictions were squared and compared with the original values – maintaining clinical relevance. Warfit-learn provides the option to perform this rooting and squaring automatically.

## 2.5 Learning algorithms used

A basic linear regression model – similar to that of the IWPC [16] – was required as a baseline for model performance. This was compared against many of the best-performing algorithms described by Liu et al. [5] and Ma et al. [12]. All algorithms were implemented using the Scikit-learn library [30] in Python 3, with the use of MLxtend [31] to implement stacked ensembles. The implementation details and key hyperparameters can be found in Table 3. Additionally, a naïve algorithm was implemented to establish a lower bound on what could be considered as legitimate modelling. This algorithm simply learns the median of the training data and then predicts that value for all subsequent cases.

## 2.6 Clinical and statistical metrics

Appropriate metrics were chosen to assess both the clinical and statistical accuracy of the models. An important clinical consideration is that INR has a therapeutic range of  $\pm 0.5$  in most patients [32] and using tighter target ranges for maintenance dosing does not achieve any therapeutic advantage [33]. The chosen metrics accounted for that, and were consistent with metrics used in related studies – allowing direct results comparisons. Moreover, the use of multiple metrics for performance (one clinically-relevant, the other statistically-relevant) is a worthwhile methodology to employ as it prevents the kind of systemic overfitting that occurs in other domains where a single metric is dominant.

1. **Mean absolute error (MAE)** is used widely across warfarin estimation studies [5, 6, 7, 34, 8, 9, 10, 11, 12].

$$MAE = \frac{\sum_{i=1}^n |y_i - \hat{y}_i|}{n} \quad (1)$$

where  $y_i$  is the actual dose and  $\hat{y}_i$  is the predicted dose.

2. **Percentage of patients with dose estimates within 20% of the actual therapeutic dose (PW20)** was used by a number of notable studies [16, 10, 9, 5, 12]. It reflects the fact that being within 0.5 points of the target INR is clinically sufficient [33].

$$PW20 = \frac{\sum_{i=1}^n f(p_i)}{n} \times 100\% \quad (2)$$

where  $f(p_i)$  for patient  $p_i$  is 1 if  $0.8y_i < \hat{y}_i < 1.2y_i$ , else 0;  $\hat{y}_i$  is the predicted dose, and  $y_i$  is the true therapeutic dose.

## 2.7 Evaluation protocol

Monte Carlo Cross-Validation (MCCV) was performed using the standard `train_test_split` function in Scikit-learn. This performs resampling *without* replacement. All models were trained and evaluated using 100 iterations of MCCV for each filtered dataset. In each iteration, the data was randomly split into a training set and a testing set in an 80/20 ratio. The model was then fit on the training set and evaluated in terms of PW20 and MAE on the testing set. The means of these results and 95% confidence intervals were calculated for each model on each dataset based on the percentile method.



Table 3: Learning algorithms replicating the work of the IWPC [16], Liu et al. [5], and Ma et al. [12] that were compared in this study.

Name	Implementation	Key hyperparameters
LR	sklearn.linear_model.LinearRegression	normalize=False, fit_intercept=True
SVR	sklearn.svm.LinearSVR	epsilon=0.0, tol=0.0001, C=1.0, loss='epsilon_insensitive'
SV	sklearn.svm.SVR	kernel='linear', cache_size=1000
RR	sklearn.linear_model.Ridge	alpha=1.0
BRT	sklearn.ensemble.GradientBoostingRegressor	loss='least squares', learning_rate=0.1, n_estimators=100
GBT	sklearn.ensemble.GradientBoostingRegressor	learning_rate=0.1, loss='lad', max_depth=4
NN	sklearn.neural_network.MLPRegressor	hidden_layer_sizes=(100, ), activation='logistic', solver='lbfgs'
Stacked_SV	mlxtend.regressor.StackingCVRegressor	regressors={GBT, SV, NN}, meta_regressor=SV, cv=5
Stacked_RR	mlxtend.regressor.StackingCVRegressor	regressors={GBT, RR, NN}, meta_regressor=RR, cv=5

The estimators used in the stacked ensembles had the same hyperparameters as the standalone estimators of the same name, as per the methodology of Ma et al. [12]. LR = (multiple) linear regression, SV(R) = support vector (regression), RR = ridge regression, BRT = boosted regression trees, GBT = gradient-boosted trees, NN = neural network. Note the slight implementation differences between SV (following the Ma et al. implementation) and SVR (following a generic implementation more similar to that of Liu et al.)

## 2.8 Warfit-learn software library

Extending past studies in warfarin dose prediction required the development of a host of data cleaning and preprocessing tools, evaluation routines, and scoring functions. To remove the need for repeating this effort in future work, this study implements a new software framework for Python-based development of warfarin dose estimation models. The aim of this library is to allow researchers to more easily replicate and extend the work of their colleagues in this domain. This can help to reduce confusion about methodology and enhance reproducibility. The library focuses on four areas, namely:

1. Seamless dataset loading, cleaning, and preprocessing.
2. Standardised implementations of scoring functions.
3. Multithreaded model training and evaluation using standardised resampling techniques.
4. Full interoperability with the Python scientific stack [35], Pandas [36], and Scikit-learn [30].

This makes it possible to replicate studies on the IWPC data with only a beginner’s level of Python experience. Supervised learning models can be constructed using the popular Scikit-learn implementations, which are handled natively by the framework. Libraries that extend the Scikit suite – such as MLxtend [31] – can be used directly with minimal configuration. Parallel processing on a user-defined number of CPU cores is used to achieve speedup on the resampling performed during evaluation and is wrapped in a user-friendly function. Functions for metrics like percentage of patients within 20% of therapeutic dose (PW20) are implemented and automatically called during evaluation. Final results are produced as Pandas dataframes, which can be exported to CSV files or LaTeX tables with a single command. The entire framework is available as open source software at <http://pypi.org/project/warfit-learn> under a GNU GPLv3 license. All results for this study were obtained using the above tools in version 0.2 of Warfit-learn.

## 2.9 Optimisation with genetic programming

TPOT [23] was used to generate high-performing machine learning pipelines through genetic programming. In each run, a preprocessed version of the dataset was given as input to TPOT and many generations of computation yielded the best performers. TPOT accepts bespoke scoring functions as its evolutionary fitness function. The functions for PW20 and MAE were trialled independently as fitness functions. We also explored a combination of these metrics. Attempting to both maximise PW20 and minimise MAE simultaneously would have transformed the task into a multi-objective optimisation. A popular approach to overcome these complexities is *scalarisation*, which numerically combines both objectives into a single optimisation function [37]. However, this requires a top-down definition of appropriate weightings for both the PW20 and MAE. To avoid introducing additional hyperparameters, we defined a new hybrid function as the ratio of the maximised (PW20) to minimised (MAE) objectives:

$$\text{hybrid} = \frac{\text{PW20}}{\text{MAE}^2} \quad (3)$$

Because PW20 was encoded as a percentage, typical values were in the range [20, 50], whilst MAE was usually below 15. We thus squared the MAE denominator so that the two values would be more similar in range, whilst also boosting the effects of variations in the error component.

Many instances of TPOT were run on a multi-core machine using different evolutionary hyperparameters. The number of generations ranged from 50 to 1000, the number of offspring from 5 to 100, and the  $k$  used in  $k$ -fold cross-validation from 5 to 30. Some runs used the raw weekly dose as the target feature and others used the square root of weekly dose. Over all the instances run on both datasets, the best performers (according to final validation scores within TPOT) were evaluated against the human-optimised algorithms described in other studies.

### 3 Results and discussion

The performance of all estimators is compared across datasets in Table 4. For easy comparison with previous studies, the mean and 95% confidence intervals are reported for both the MAE and PW20 metrics. In order to more clearly illustrate and compare the distributions of the results, box plots are presented in Figure 2.

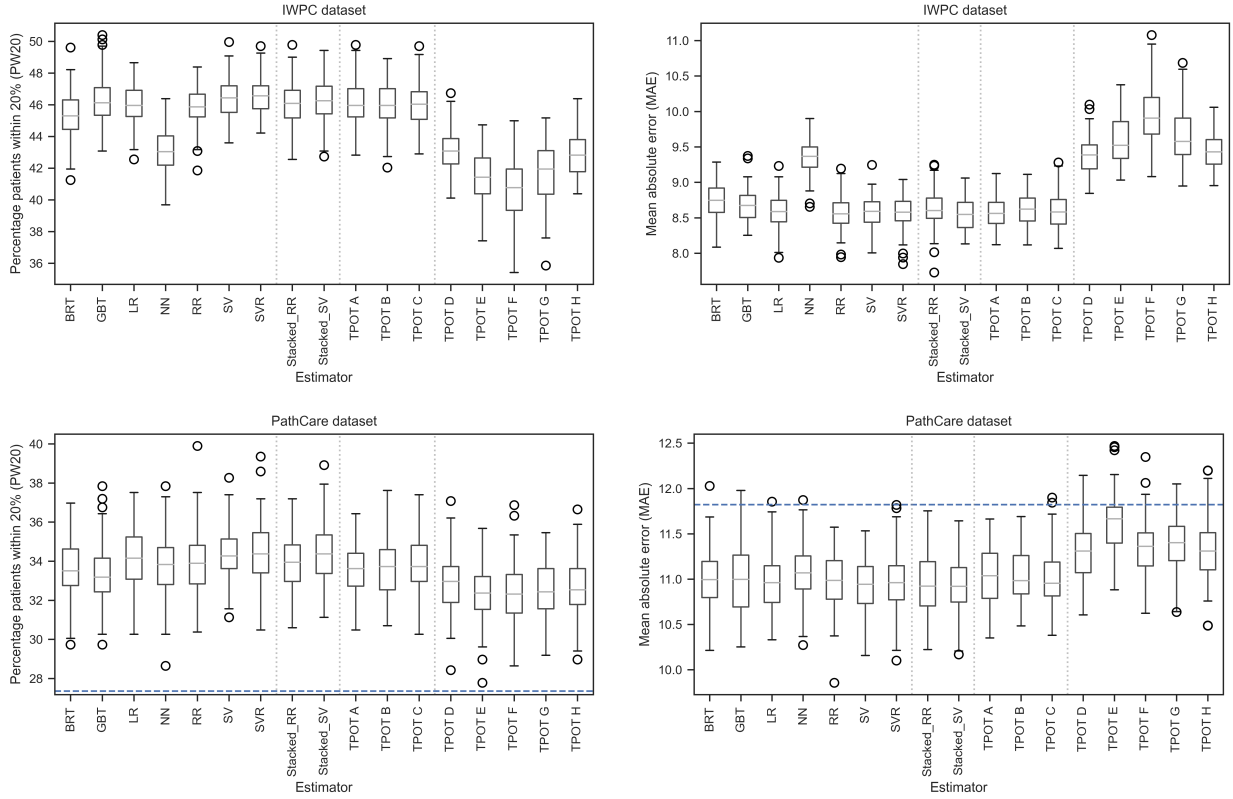


Figure 2: Box plots showing the distributions of all estimators after 100 rounds of MCCV evaluation on both the IWPC (top row) and PathCare (bottom row) datasets. Each box marks the first and third quartiles with a line at the median. The whiskers extend to the full range of the data. Results outside 1.5 times the interquartile range are considered outliers and marked with circles. The dotted vertical lines delineate the traditional algorithms, the stacked ensembles from Ma et al. [12], the IWPC-optimised TPOT architectures, and the PathCare-optimised architectures. The dashed blue lines on the PathCare plots (bottom row) indicate the mean performance of a naïve estimator. The axes are automatically scaled to best illustrate the variation between estimators. BRT = boosted regression trees, GBT = gradient-boosted trees, LR = (multiple) linear regression, NN = neural network, RR = ridge regression, SV(R) = support vector (regression). The naïve estimator always predicted the median of the training data and was used as a lower bound on what is considered genuine modelling. See the implementation details for each estimator in Table 3. Details of the pipelines generated by TPOT can be found with the published source code for this project.

Table 4: Results of all estimators after 100 rounds of MCCV evaluation on both the IWPC and novel (PathCare) datasets, with the best performers in each group emphasised.

Estimator	IWPC Dataset			PathCare Dataset		
	PW20 (95% CI)	MAE (95% CI)	PW20 (95% CI)	MAE (95% CI)	PW20 (95% CI)	MAE (95% CI)
BRT	45.32 (42.47–47.70)	8.75 (8.35–9.18)	33.61 (31.19–36.27)	10.99 (10.40–11.52)		
GBT	46.31 (43.73–49.74)	8.67 (8.32–9.07)	33.35 (30.42–36.60)	10.99 (10.37–11.60)		
LR	45.99 (43.38–48.22)	8.59 (8.11–9.05)	34.08 (30.81–36.82)	10.98 (10.39–11.67)		
NN	43.12 (40.69–45.87)	9.36 (8.90–9.80)	33.79 (31.19–37.14)	11.07 (10.56–11.67)		
RR	45.96 (43.33–48.30)	8.57 (8.16–9.03)	33.99 (31.19–36.92)	10.98 (10.45–11.53)		
SV	46.41 (44.07–48.74)	8.58 (8.12–8.96)	34.40 (31.88–36.81)	10.93 (10.36–11.45)		
<b>SVR</b>	<b>46.56 (44.60–49.13)</b>	<b>8.58 (8.05–9.00)</b>	<b>34.46 (31.77–37.14)</b>	<b>10.95 (10.33–11.60)</b>		
Stacked_RR	46.01 (43.29–48.30)	8.63 (8.19–9.14)	33.83 (30.86–36.55)	10.95 (10.39–11.54)		
<b>Stacked_SV</b>	<b>46.32 (43.81–48.48)</b>	<b>8.55 (8.21–8.99)</b>	<b>34.45 (31.78–37.62)</b>	<b>10.93 (10.26–11.57)</b>		
TPOT A*	46.03 (43.03–48.26)	8.61 (8.17–9.06)	33.65 (31.29–36.54)	11.05 (10.58–11.66)		
TPOT B*	46.03 (43.56–48.83)	8.59 (8.16–9.06)	33.91 (31.83–36.49)	11.01 (10.44–11.69)		
<b>TPOT C*</b>	<b>46.08 (43.73–48.53)</b>	<b>8.56 (8.14–9.02)</b>	<b>33.53 (31.03–35.85)</b>	<b>11.03 (10.49–11.53)</b>		
TPOT D <sup>‡</sup>	43.09 (40.55–45.88)	9.38 (8.93–9.88)	32.86 (30.32–35.75)	11.29 (10.75–11.98)		
TPOT E <sup>‡</sup>	41.44 (38.19–44.30)	9.59 (9.07–10.17)	32.32 (29.62–34.99)	11.63 (11.03–12.29)		
TPOT F <sup>‡</sup>	40.54 (36.38–43.78)	9.96 (9.28–10.86)	32.37 (29.95–35.20)	11.35 (10.74–11.93)		
TPOT G <sup>‡</sup>	41.75 (38.68–44.92)	9.64 (9.05–10.58)	32.53 (30.21–35.25)	11.39 (10.83–11.95)		
TPOT H <sup>‡</sup>	42.82 (40.64–45.60)	9.43 (8.99–9.89)	32.67 (29.83–35.41)	11.33 (10.80–12.07)		
Naïve	34.10 (33.85–34.35)	12.45 (12.39–12.50)	27.35 (27.10–27.61)	11.82 (11.75–11.89)		

MAE = mean absolute error. PW20 = percentage of patients within 20% of therapeutic dose. Values are provided as means with 95% confidence intervals. All estimators utilised the same parameter set for the IWPC data, as defined by the seminal IWPC study [16]. \* = Evolved on IWPC dataset. † = Evolved on PathCare dataset using raw weekly dose. ‡ = Evolved on PathCare dataset using square root of weekly dose. BRT = boosted regression trees, GBT = gradient-boosted trees, LR = (multiple) linear regression, NN = neural network, RR = ridge regression, SV(R) = support vector (regression). The native estimator always predicted the median of the training data and was used as a lower bound on what is considered genuine modelling. See the implementation details for each estimator in Table 3. Details of the pipelines generated by TPOT can be found with the published source code for this project.

### 3.1 Validation of methodology

The methodology of this study was validated by referencing the results of previous work on warfarin dose estimation. The original IWPC study reported an MAE of 8.5 mg/week (95% CI: 8.1–8.6) for their linear regression (LR) model [16]. However, they processed the cohort differently.

The study by Ma et al. evaluated models on an identical dataset to our study, reporting an MAE of 8.53 (95% CI: 8.08–8.99) for their LR model [12]. By comparison, the MAE for the LR model in this study was 8.59 mg/week (95% CI: 8.11–9.05). This indicates that our preprocessing and evaluation methodology were valid. It also justifies the original IWPC finding that, when adequate preprocessing is performed, linear regression produces outstanding results on the IWPC dataset [16]. This, coupled with its easy interpretability, makes LR hard to rule out as the default warfarin estimation model. However, as datasets grow in size and variety, LR can fail to model some relationships that more complex approaches succeed in capturing. It also tends to be less robust to outliers than ensemble approaches [38], which could challenge its clinical applicability.

The shapes of the performance distributions in Figure 2 are quite symmetric and consistent with normal distributions, which affirms the use of confidence intervals based on the percentile method. Despite some large differences in average performance over the 100 resamplings, the distributions still overlap considerably, highlighting how much performance can vary depending on the specific patients in the training and testing sets. With few exceptions, the ranges of PW20 and MAE scores for even the most consistent algorithms were considerably greater than the differences between the mean scores, suggesting that both datasets have a high degree of inter-patient variation that is unaccounted for in the available parameters.

The purpose of mirroring our methodology on the PathCare dataset was to evaluate how well the estimators generalise to a different clinical context and parameter set. The relationships of the IWPC results are certainly reflected in the PathCare results, both in terms of relative scores (Table 4) and variance (Figure 2). However, they are much less pronounced.

To verify that effective modelling took place, we made use of a naïve estimator that simply predicted the median of its training set for each resampling. On the IWPC dataset, it had a PW20 of 34.10 (95% CI: 33.85–34.35) and an MAE of 12.45 (95% CI: 12.39–12.50), making it the lowest performer by a considerable amount. On the PathCare dataset, it achieved a PW20 of 27.35 (95% CI: 27.10–27.61) and an MAE of 11.82 (95% CI: 11.75–11.89). The other estimators only marginally outperformed the naïve baseline on the PathCare dataset in terms of MAE, but well outperformed it in terms of PW20 – the clinically-relevant metric. This affirms that even a limited number of clinical parameters are sufficient for the estimators to learn some relationships from the data and make meaningful predictions. For the PathCare dataset, the mean performance of the naïve baseline is overlaid in the box plots of Figure 2 for visual reference.

### 3.2 Performance of traditional techniques

Support vector regression (SV and SVR) was a top-performing algorithm, whilst the neural network (NN) was one of the worst performing algorithms. The relative performance of the NN and SVR models was similar to that reported by Liu et al. [5]. Their NN (MAE 9.82, PW20 41.27) was by far their worst performer, whilst their SVR (MAE 8.96, PW20 45.88) was one of their best performers. Similarly, this study’s NN (MAE 9.36, PW20 43.12) was the worst-performing standalone algorithm by a considerable margin, whilst our SVR (MAE 8.58, PW20 46.56) was the best.

Unlike most statistical models, neural networks are considered to be good at learning feature representations automatically, which is especially useful when there are many features. The IWPC dataset has very few parameters compared to typical deep learning datasets. Moreover, these parameters are heavily-engineered during the preprocessing stage, which likely minimises any effects of the representation learning. It is also widely held that NN performance scales with the size of the dataset and the number of hidden layers [39]. Because the IWPC dataset contains very few records

(by deep learning standards) using more than one hidden layer has little-to-no effect. These factors may explain why the NN consistently failed to perform on both datasets.

Our results replicate the findings of Ma et al. [12] for LR, SV, and ridge regression (RR) algorithms, with very similar confidence intervals. But, despite replicating the implementations outlined in their methodology, our results for NN, gradient boosting trees (GBT), Stacked SV, and Stacked RR were notably dissimilar. In the case of the ensemble estimators, this was likely due to their decision to use a much larger parameter set on their stacked implementations than they fed to other models. In contrast, our study made use of the same parameter set across all estimators (as standardised by notable previous works [16, 5]). Another factor may have been the difference in GBT implementations (Scikit-learn vs. LightGBM) between studies. This discrepancy once again affirms the need for a standardised set of tools for replicating and evaluating warfarin dosing models.

Despite the difference in absolute scores between studies, the stacked ensembles were still among the top performers on the IWPC data, with the Stacked SV in particular achieving an MAE of 8.55 and a PW20 of 46.32. This supports the use of the stacked generalisation approach to improve warfarin dose estimation. However, with a much smaller effect size than Ma et al. reported, it may be difficult to justify the use of these techniques over the more-explainable linear regression model in clinical deployment.

Despite the absence of pharmacogenetic data and many clinical parameters in the PathCare dataset, the traditional algorithms performed reasonably well. LR (MAE 10.98, PW20 34.08) and SV (MAE 10.93, PW20 34.40) topped the list, along with the Stacked SV ensemble (MAE 10.93, PW20 34.45). However, there was very little difference in performance on the PathCare dataset amongst the traditional algorithms. The range of the distributions were also much wider on the PathCare dataset than the IWPC dataset (see Figure 2). The fact that all the estimators still outperformed the naïve baseline indicates that there were indeed predictive relationships to be found in the PathCare dataset, but it seems that the precise choice of algorithm was of much less importance with so few parameters available. Algorithms with higher bias, like LR and SV, performed marginally better, whereas algorithms with lower bias, like the ensembles, appear to have generalised less well.

### 3.3 Performance of evolved architectures

The architectures developed using a genetic programming approach (TPOT) had extremely varied results depending on which dataset they were evolved on. TPOT C was an estimator pipeline evolved on the IWPC dataset that produced the best results of the evolved models (MAE 8.56, PW20 46.08). It was only marginally outperformed by the SV and Stacked SV approaches (with highly similar CIs), but outperformed the Stacked RR estimator on the IWPC dataset. The structure and hyperparameters of the estimator are illustrated in Figure 3.

Architectures evolved on the IWPC dataset (A, B, and C) performed well on the IWPC dataset, with almost identical distributions. This is remarkable when considering that they were generated with no human tuning. It is also notable that three different evolutionary histories converged to have such consistent performance. This result concurs with previous findings [22, 23] that autoML can attain results comparable to those of models engineered by experts with domain knowledge.

Generally, the TPOT architectures that performed well on the IWPC data also performed well on the PathCare data. This could indicate that the IWPC dataset captures sufficient information about the parameter-target relationships to generalise, but may also be a result of limitations in the PathCare dataset.

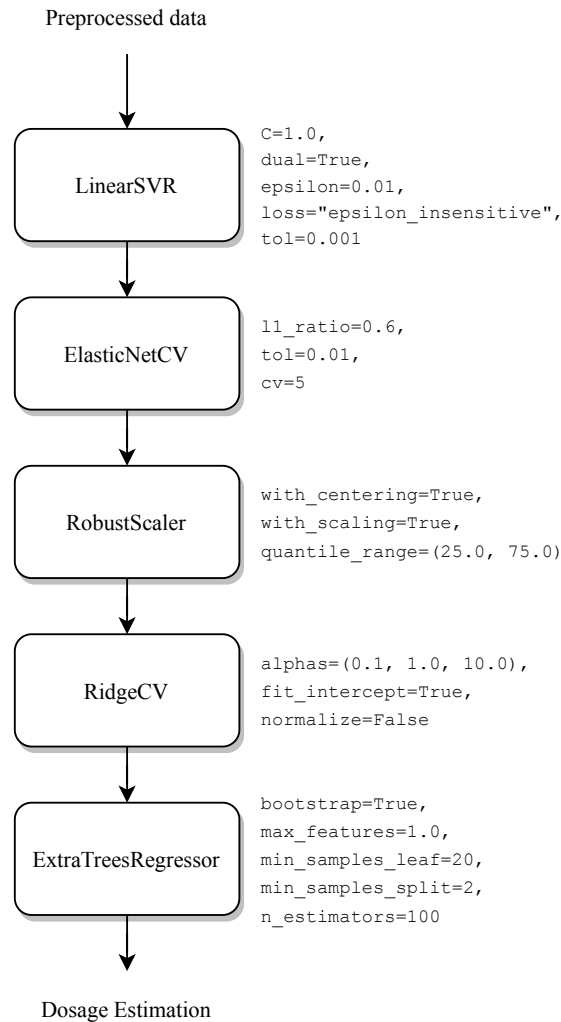


Figure 3: Illustration of the highest-performing architecture and hyperparameters evolved on the IWPC dataset using TPOT. The Python implementation of this architecture can be found along with the project source code.

The architectures evolved on the PathCare dataset (D through H) performed significantly worse than those evolved on the IWPC dataset (A through C) regardless of which dataset they were trained and evaluated on. Moreover, the distributions of results were much wider, indicating low generalisation. These performance differences are best illustrated by the jarring shift in distributions seen in the top row of Figure 2. These results suggest that the PathCare dataset has a lower ratio of signal to noise compared to the IWPC dataset and support prior claims [40, 41, 42] that pharmacogenetic parameters yield more generalised models.

All but one of the architectures evolved with TPOT utilised stacking at some point in their pipelines. This may reflect a general tendency within the TPOT implementation but, when considered with the high performance of Ma et al.’s stacked ensembles, may indicate the robustness of stacking for warfarin dose estimation. Only TPOT C (the best overall) and TPOT E included SVR in their pipelines. This is strange considering the high level of performance seen with SVR across both datasets. TPOTs B and C (high performers on the IWPC dataset) made use of trees (with boosting) and/or lasso regression. TPOTs D-H (evolved on PathCare) all relied on a KNN regressor somewhere in their pipeline, but performed poorly across datasets.

### 3.4 Limitations

The IWPC dataset was compiled by 22 research groups, each with different protocols and equipment. This resulted in noisy data and some missing values, which lower the predictive accuracy of models. It is also known that the impact of *CYP2C9* and *VKORC1* genes varies across races [5]. Current pharmacogenetic data is mostly derived from White and Asian OMB racial groups in developed nations (where genetic testing is available), so it is likely that current pharmacogenetic implementations impart bias upon models.

The PathCare dataset did not include clinical information such as height, weight, and race – all of which have been shown to improve dosing accuracy. It would, therefore, be imprudent to extrapolate very much from the performance differences of models on this dataset. Nevertheless, it served as a useful validation that top-performing algorithms can generalise to novel clinical contexts with different parameter sets. Because of these different parameters, it was not possible to directly train models on one dataset and evaluate their performance on another. It was, however, possible to evolve the architectures and hyperparameters on one dataset and then train and evaluate models on the other dataset – which was done and is presented in Table 4 and Figure 2.

There are also technical limitations in the performance of the genetic programming tools. As Olson et al. point out themselves, TPOT becomes quite slow as the size of a dataset increases. Even on the tiny (by machine learning standards) IWPC dataset, a single run could take days to complete. There are plans to integrate other autoML libraries and heuristics to seed the populations with strong pipelines and reduce the number of computations to convergence [23]. Additionally, the use of highly-parallel hardware, such as graphics processing units (GPUs), offers enormous potential for speedup in execution. Because TPOT is built atop the popular, open-source Scikit-learn framework, it benefits from modularity and standardisation. Both the *H2O4GPU* project and Nvidia’s *cuML* suite mirror the Scikit-learn interface, but move computations to the GPU. At the time of writing, there is at least one ongoing project to bring this functionality into TPOT.

Unlike linear regression, the other models do not produce human-interpretable formulae for clinical warfarin dosing. However, trained models could be used to run inference on real patient data and produce dosage recommendations to help guide clinicians. Further work is needed to produce robust models that can be deployed as decision-support systems in a clinical context.

## 4 Conclusions

This study presented the Warfit-learn software framework, which implements the best published techniques for warfarin dosage modelling. The framework was validated by replicating the results of several prior studies on the same IWPC dataset. Warfit-learn’s parallel evaluation methodology was utilised to compare the performance of several estimators across both the IWPC dataset and a novel dataset of South African patients provided by PathCare. The estimators included traditional learning algorithms that have shown promise in warfarin dosing, two stacked generalisation ensembles described in previous work, and a collection of learning pipelines evolved using a genetic programming approach to architecture and hyperparameter optimisation.

As in previous studies, linear regression was shown to be a reliable technique. Support vector regression was the best performing traditional algorithm, whilst neural networks performed poorly. Recent ensemble approaches – namely stacked generalisation – were shown to be effective across datasets, but the effect size was not as large as previously reported.

This study found that genetic programming was effective at evolving robust architectures and hyperparameters for warfarin dosage modelling. All three of the estimators evolved on the IWPC dataset were amongst the top-performing models, performing on par with the best expert-designed ensembles. Inspection of the architectures revealed a preference for stacked ensembles. If these results are not unique to warfarin dosing, they suggest that automated machine learning may be a promising method for the future of software-assisted dosing. These findings should be explored on other



dosage estimation problems in future work. Other advances in this domain can be made by compiling comprehensive warfarin dosing datasets from more diverse cohorts, and using them to train robust models for safe clinical deployment.

In addition to the performance comparisons and evolutionary optimisation techniques, this paper presented a much-needed software framework for evaluating new techniques on the IWPC dataset. Warfit-learn eliminates the need for future researchers to re-create tools, and maintains consistency by providing a standardised testbed. It is our hope that Warfit-learn will make future studies significantly easier and more reproducible.

## References

- [1] Daniel E. Jonas and Howard L. McLeod. Genetic and clinical factors relating to warfarin dosing. *Trends in Pharmacological Sciences*, 30(7):375–386, 2009.
- [2] P S Wells, A M Holbrook, N R Crowther, and J Hirsh. Interactions of warfarin with drugs and food, nov 1994.
- [3] T B L Kirkwood and S M Lewis. Requirements for thromboplastins and plasma used to control oral anticoagulant therapy. *WHO Tech Rep Ser*, 687:81–99, 1983.
- [4] L. Poller. International Normalized Ratios (INR): The first 20 years. *Journal of Thrombosis and Haemostasis*, 2(6):849–860, 2004.
- [5] Rong Liu, Xi Li, Wei Zhang, and Hong Hao Zhou. Comparison of nine statistical model based warfarin pharmacogenetic dosing algorithms using the racially diverse international warfarin pharmacogenetic consortium cohort database. *PLoS ONE*, 10(8), 2015.
- [6] Yu Liu, Jie Yang, Qiang Xu, Bin Xu, Lei Gao, Yuxiao Zhang, Yan Zhang, Hongjuan Wang, Caiyi Lu, Yusheng Zhao, and Tong Yin. Comparative performance of warfarin pharmacogenetic algorithms in Chinese patients. *Thrombosis Research*, 130(3):435–440, 2012.
- [7] Ashkan Sharabiani, Adam Bress, Elnaz Douzali, and Houshang Darabi. Revisiting warfarin dosing using machine learning techniques. *Computational and Mathematical Methods in Medicine*, 2015, 2015.
- [8] Ya Han Hu, Fan Wu, Chia Lun Lo, and Chun Tien Tai. Predicting warfarin dosage from clinical data: A supervised learning approach. *Artificial Intelligence in Medicine*, 56(1):27–34, 2012.
- [9] Qin Zhou, Joey Kwong, Jie Chen, Wenzhe Qin, Jin Chen, and Li Dong. Use of artificial neural network to predict warfarin individualized dosage regime in Chinese patients receiving low-intensity anticoagulation after heart valve replacement. *International Journal of Cardiology*, 176(3):1462–1464, oct 2014.
- [10] Enzo Grossi, Gian Marco Podda, Mariateresa Pugliano, Silvia Gabba, Annalisa Verri, Giovanni Carpani, Massimo Buscema, Giovanni Casazza, and Marco Cattaneo. Prediction of optimal warfarin maintenance dose using advanced artificial neural networks. *Pharmacogenomics*, 15(1):29–37, jan 2014.
- [11] Ashkan Sharabiani, Houshang Darabi, Adam Bress, Larisa Cavallari, Edith Nutescu, and Katarzyna Drozda. Machine learning based prediction of warfarin optimal dosing for African American patients. In *Automation Science and Engineering (CASE), 2013 IEEE International Conference on*, pages 623–628, 2013.
- [12] Zhiyuan Ma, Ping Wang, Zehui Gao, Ruobing Wang, and Koroush Khalighi. Ensemble of machine learning algorithms using the stacked generalization approach to estimate the warfarin dose. *Plos One*, 13(10):e0205872, 2018.
- [13] Leon Poller, M. Keown, S. Ibrahim, G. Lowe, M. Moia, A. G. Turpie, C. Roberts, A. M.H.P. Van Den Besselaar, F. J.M. Van Der Meer, A. Tripodi, G. Palareti, C. Shiach, S. Bryan, M. Samama, M. Burgess-Wilson, A. Heagerty, P. Maccallum, D. Wright, and J. Jespersen. An international multicenter randomized study of computer-assisted oral anticoagulant dosage vs. medical staff dosage. *Journal of Thrombosis and Haemostasis*, 6(6):935–943, jun 2008.

- [14] S. McDonald, C. Xydeas, and P. Angelov. A retrospective comparative study of three data modelling techniques in anticoagulation therapy. *BioMedical Engineering and Informatics: New Development and the Future - Proceedings of the 1st International Conference on BioMedical Engineering and Informatics, BMEI 2008*, 1:219–225, 2008.
- [15] Y. K. Kim, R. Nieuwlaat, S. J. Connolly, S. Schulman, K. Meijer, N. Raju, S. Kaatz, and J. W. Eikelboom. Effect of a simple two-step warfarin dosing algorithm on anticoagulant control as measured by time in therapeutic range: A pilot study. *Journal of Thrombosis and Haemostasis*, 8(1):101–106, 2010.
- [16] B.F. Klein, T.E., Altman, R.B., Eriksson. Estimation of the Warfarin Dose with Clinical and Pharmacogenetic Data. *Archives of Internal Medicine*, 360(8):753–764, 2009.
- [17] M Whirl-Carrillo, EM McDonogh, J Herbet, L Gong, K Sangkuhl, C Thotn, R Altman, and E Klein. Pharmacogenomics Knowledge for Personalized Medicine. *Clinical Pharmacology and Therapeutics*, 92(4):414–417, 2012.
- [18] B.F. Klein, T.E., Altman, R.B., Eriksson. Supplement to: The International Warfarin Pharmacogenetics Consortium. Estimation of the warfarin dose with clinical and pharmacogenetic data. *Journal of Medicine (Cincinnati)*, 2006.
- [19] Erdal Cosgun, Nita A. Limdi, and Christine W. Duarte. High-dimensional pharmacogenetic prediction of a continuous trait using machine learning techniques with application to warfarin dose prediction in African Americans. *Bioinformatics*, 27(10):1384–1389, 2011.
- [20] Frank Hutter, Jörg Lücke, and Lars Schmidt-Thieme. Beyond manual tuning of hyperparameters. *KI-Künstliche Intelligenz*, 29(4):329–337, 2015.
- [21] James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13(Feb):281–305, 2012.
- [22] Jason Zutty, Daniel Long, Heyward Adams, Gisele Bennett, and Christina Baxter. Multiple objective vector-based genetic programming using human-derived primitives. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, pages 1127–1134. ACM, 2015.
- [23] Randal S Olson, Nathan Bartley, Ryan J Urbanowicz, and Jason H Moore. Evaluation of a Tree-based Pipeline Optimization Tool for Automating Data Science. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016, GECCO '16*, pages 485–492, New York, NY, USA, 2016. ACM.
- [24] Gregory S Hornby, Jason D Lohn, and Derek S Linden. Computer-automated evolution of an X-band antenna for NASA’s space technology 5 mission. *Evolutionary computation*, 19(1):1–23, 2011.
- [25] Erik M Fredericks and Betty H C Cheng. Exploring automated software composition with genetic programming. In *Proceedings of the 15th annual conference companion on Genetic and evolutionary computation*, pages 1733–1734. ACM, 2013.
- [26] Stephanie Forrest, ThanhVu Nguyen, Westley Weimer, and Claire Le Goues. A genetic programming approach to automated software repair. In *Proceedings of the 11th Annual conference on Genetic and evolutionary computation*, pages 947–954. ACM, 2009.
- [27] Lee Spector, David M Clark, Ian Lindsay, Bradford Barr, and Jon Klein. Genetic programming for finite algebras. In *Proceedings of the 10th annual conference on Genetic and evolutionary computation*, pages 1291–1298. ACM, 2008.
- [28] Wolfgang Banzhaf, Peter Nordin, Robert E Keller, and Frank D Francone. *Genetic programming: an introduction*, volume 1. Morgan Kaufmann San Francisco, 1998.
- [29] Félix-Antoine Fortin, François-Michel De Rainville, Marc-André Gardner, Marc Parizeau, and Christian Gagné. DEAP: Evolutionary algorithms made easy. *Journal of Machine Learning Research*, 13(Jul):2171–2175, 2012.
- [30] F Pedregosa, G Varoquaux, A Gramfort, V Michel, B Thirion, O Grisel, M Blondel, P Prettenhofer, R Weiss, V Dubourg, J Vanderplas, A Passos, D Cournapeau, M Brucher, M Perrot, and E Duchesnay. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

- [31] Sebastian Raschka. MLxtend: Providing machine learning and data science utilities and extensions to Python's scientific computing stack. *The Journal of Open Source Software*, 3(24), apr 2018.
- [32] G. W. Albers, J. E. Dalen, A. Laupacis, W. J. Manning, P. Petersen, and D. E. Singer. Antithrombotic therapy in atrial fibrillation. *Chest*, 119(1 SUPPL.), 2001.
- [33] D. J. Meier, S. Seva, and William P. Fay. A comparison of anticoagulation results of patients managed with narrow vs. standard international normalized ratio target ranges [6]. *Journal of Thrombosis and Haemostasis*, 5(6):1332–1334, 2007.
- [34] S. L. Tan, Z. Li, G. B. Song, L. M. Liu, W. Zhang, J. Peng, T. Zhang, F. F. Jia, G. Zhou, H. H. Zhou, and X. M. Zhou. Development and comparison of a new personalized warfarin stable dose prediction algorithm in Chinese patients undergoing heart valve replacement. *Pharmazie*, 67(11):930–937, 2012.
- [35] K. Jarrod Millman and Michael Aivazis. Python for Scientists and Engineers. *Computing in Science & Engineering*, 13(2):9–12, mar 2011.
- [36] Wes McKinney. *Data Structures for Statistical Computing in Python*, 2010.
- [37] Agoston E Eiben and James E Smith. *Introduction to evolutionary computing*. Springer, 2015.
- [38] Christopher M Bishop. *Pattern recognition and machine learning*. springer, 2006.
- [39] Anna Choromanska, Mikael Henaff, Michael Mathieu, Gérard Ben Arous, and Yann LeCun. The Loss Surfaces of Multilayer Networks. *arXiv*, 38, 2014.
- [40] Munir Pirmohamed, Girvan Burnside, Niclas Eriksson, Andrea L. Jorgensen, Cheng Hock Toh, Toby Nicholson, Patrick Kesteven, Christina Christersson, Bengt Wahlström, Christina Stafberg, J. Eunice Zhang, Julian B. Leathart, Hugo Kohnke, Anke H. Maitland-van der Zee, Paula R. Williamson, Ann K. Daly, Peter Avery, Farhad Kamali, and Mia Wadelius. A Randomized Trial of Genotype-Guided Dosing of Warfarin. *New England Journal of Medicine*, 369(24):2294–2303, 2013.
- [41] Elizabeth A. Sconce, Tayyaba I. Khan, Hilary A. Wynne, Peter Avery, Louise Monkhouse, Barry P. King, Peter Wood, Patrick Kesteven, Ann K. Daly, and Farhad Kamali. The impact of CYP2C9 and VKORC1 genetic polymorphism and patient characteristics upon warfarin dose requirements: Proposal for a new dosing regimen. *Blood*, 106(7):2329–2333, 2005.
- [42] Mia Wadelius, Leslie Y. Chen, Niclas Eriksson, Suzannah Bumpstead, Jilur Ghori, Claes Wadelius, David Bentley, Ralph McGinnis, and Panos Deloukas. Association of warfarin dose with genes involved in its action and metabolism. *Human Genetics*, 121(1):23–34, 2007.