

Feature importance in machine learning models: A fuzzy information fusion approach



Divish Rengasamy^{b,1}, Jimiama M. Mase^{a,*}, Aayush Kumar^a, Benjamin Rothwell^c, Mercedes Torres Torres^a, Morgan R. Alexander^d, David A. Winkler^{e,f,g}, Graziela P. Figueredo^a

^a School of Computer Science, University of Nottingham, UK

^b Institute for Aerospace Technology, The University of Nottingham, UK

^c Gas Turbine and Transmissions Research Centre, The University of Nottingham, UK

^d School of Pharmacy, The University of Nottingham, UK

^e Monash Institute for Pharmaceutical Sciences, Monash University, Australia

^f La Trobe Institute for Molecular Science, La Trobe University, Australia

^g School of Pharmacy, University of Nottingham, UK

ARTICLE INFO

Article history:

Received 23 May 2022

Revised 29 July 2022

Accepted 4 September 2022

Available online 13 September 2022

Keywords:

Feature importance

Fuzzy systems

Information fusion

Interpretability

Machine learning

Responsible AI

ABSTRACT

With the widespread use of machine learning to support decision-making, it is increasingly important to verify and understand the reasons why a particular output is produced. Although post-training feature importance approaches assist this interpretation, there is an overall lack of consensus regarding how feature importance should be quantified, making explanations of model predictions unreliable. In addition, many of these explanations depend on the specific machine learning approach employed and on the subset of data used when calculating feature importance. A possible solution to improve the reliability of explanations is to combine results from multiple feature importance quantifiers from different machine learning approaches coupled with re-sampling. Current state-of-the-art ensemble feature importance fusion uses crisp techniques to fuse results from different approaches. There is, however, significant loss of information as these approaches are not context-aware and reduce several quantifiers to a single crisp output. More importantly, their representation of “importance” as coefficients may be difficult to comprehend by end-users and decision makers. Here we show how the use of fuzzy data fusion methods can overcome some of the important limitations of crisp fusion methods by making the importance of features easily understandable.

© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The growing application of machine learning (ML) to a rapidly increasing number of important applications has resulted in sophisticated and effective solutions with little to no human intervention. As the complexity of these ML models increases, however, understanding why and how decisions are made is important yet challenging. For safety-critical systems, understanding how ML outputs are generated is especially important for model verification, regulatory compliance [1], elucidation of ethical concerns, trustworthiness [2], model diagnostics [3–7] and design of the next generation of intelligent machines, amongst many other applications. One way to explain ML models outputs is by post-training

analysis, such as feature importance (FI) analysis. FI estimates the contribution of each feature (independent parameter) to the model output [8,4]. Using FI to understand how decisions are made aids the establishment of true causality between important data attributes and outcomes in model inference [6,9]. In interdisciplinary contexts, where ML researchers work with other domain experts, robust identification of important variables and their influence on ML model outputs facilitates expert validation of the domain-specific processes being modelled [10]. The elucidation of mechanistic processes linking the features to the outputs can validate results and lead to confidence in the model [11–13], acceptance, and reassurance of a fair intelligent tool being implemented [14,15]; or reveal biases in data or model training and establish new requirements for model refinement.

The wide availability of ML algorithms and diversity of FI techniques complicates model selection, FI approaches used, reliability of interpretation of results from a given ML method coupled to a FI

* Corresponding author.

E-mail address: Jimiama.MafeniMase1@nottingham.ac.uk (J.M. Mase).

¹ First and second authors have equally contributed to the work.

approach, and data representativeness. Different ML models may generate different FI values due to variations in their learning algorithms. Similarly, different FI techniques may produce different importances with the same ML algorithms or training data. For example, models that can adequately map the input to output relationships using linear functions can generate clear FIs, while for relationships that are significantly nonlinear, FI can be a local, context-dependent property of the response surface. To add to the complication, some FI methods are model-agnostic, while others are model-specific.

To address these uncertainties in FI, Rengasamy et al. [4] proposed an ensemble feature importance (EFI) method. Here, multiple FI approaches are applied to a set of ML models and their crisp importance values are combined to produce a final importance for each feature. This resulted in more robust and accurate analysis (15% less FI error) of synthetic data sets compared to the use of a single ML and FI. However, their study had sampling limitations (train-test split) that prevented a comprehensive exploration of the feature space and feature causality. Furthermore, they ignored the fact that the ensemble FI values contain significantly more information than that captured by crisp values. Most importantly, representation of FI as crisp coefficients can be difficult to understand or interpret for decision making.

To overcome those limitations, we propose a novel Fuzzy information fusion method known as FEFI (Fuzzy Ensemble Feature Importance) that captures and models the variance of different ML methods and FI techniques used to generate FI and data space representation. It transforms and combines the crisp importance values for each feature from different ML models using Fuzzy Sets (FSs) [16] into 'low', 'moderate' and 'high' categories of importance. FEFI outputs three sets of membership functions (MFs) [16] determining:

1. for each ML approach, what is the range of low, moderate and high importance in the feature set after training;
2. for each feature, its low, moderate and high importance relative to the other features in the data for each ML approach;
3. for each feature, its low, moderate, high importance after training and after combining with importances from multiple ML approaches.

The sets of MFs are obtained from the FI coefficients generated by applying multiple (ML, FI) combinations to the dataset using resampling (e.g., by cross-validation). This allows the variance between the ML models, FI techniques and data to be modelled as FSs, where 'low', 'moderate' and 'high' are the linguistic labels representing the relative importance of a feature. The Wang-Mendel [17] approach is employed to learn the rules that map the set of FI values to the final description of importance. Lastly, the Mamdani inference approach [18] is used to combine the FI terms generated by the MFs, using rules generated by Wang-Mendel, to produce a final description of feature importance that is easy to interpret for decision-making.

2. Background

2.1. Problem statement

Interpretation of feature importance in linear models is relatively straightforward. The regression coefficients for each feature in a multiple linear regression model (or generalised linear model) describe the magnitude and sign of their effect on the dependent variable. The importance of each feature is a global and invariant property across the model response surface (which consists of a set of linear hyperplanes). However, for ML models that exhibit sig-

nificant non-linearity, the response surface is curved. This complicates the interpretation of feature importance because these are now local rather than global properties.

Estimating the importance of features in ML predictive analytics is currently very unreliable. Different ML models, FI techniques, and subsets of data generate different importance coefficients, often with diverse magnitudes, for the same features [4]. These uncertainties in identifying the contribution of features to ML outputs are due to:

1. For nonlinear response surfaces (the multidimensional surfaces encoded by trained ML models), the FI depends on which part of the response surface is interrogated. The variation in feature gradients from one position to another is a result of the nonlinear behaviour of real-world systems. This variation means that for nonlinear models, the FI is not fixed, but is context dependent. Sensible choices for the part of the response surfaces at which to evaluate FI are typical local or global minima or maxima in the response variable, depending on what type of optimum is relevant.
2. The variance due to the characteristics of the learning process of ML models. For example, random forests [19] combine the outputs of decision trees at the end of the training process while gradient boosting machines [20] combine them at the start of the process. Although random forest and gradient boosting machines are tree-based methods that use impurity-based methods to calculate feature importance, the final feature importances for two ML methods will differ due to their variations in training algorithms.
3. Differences in how FI is calculated and interpreted by FI techniques. For instance, certain approaches, such as permutation importance (PI) [21] investigate how each feature affects the model response individually. Here, PI shuffles the instance values for a particular feature, while maintaining the original values for the remaining features. The feature with shuffled instances is considered important if the model's performance decreases. Other approaches, however, investigate the importance of the feature both individually and in synergy with other subsets of features. For example, SHapley Additive exPlanations (SHAP) [22] calculates the contribution of features for every possible combination of the feature set investigated.
4. FI coefficients in general being calculated as the average (or weighted average) of the importance of a feature within a data sample. Information about the context (or data subspace) in which a feature has higher or lower importance during training is lost.

2.2. Related work

Earlier ensemble FI strategies for ML models focused on combining the importance of features from multiple decision trees. For example, Breiman [19] used the Gini impurity metric across decision trees to calculate feature importance. Subsequently, De Bock et al. [23] extended the idea of feature-importance fusion from multiple weak learners to generalised additive models (GAM). Here, each weak learner undergoes permutation importance (PI) to calculate FI and the learners results are averaged. Zhai and Chen [24] improved ensemble FI by using multiple ML models and gains in gini importance defining the final FI. Rengasamy et al. [4], proposed a model agnostic ensemble FI framework to improve FI quantification using multiple models and multiple FI calculation methods. They studied several crisp fusion metrics such as mean, median, majority vote, rank correlation, combination with majority vote, and modified Thompson tau test. Majority voting produced the best overall results for several synthetic datasets. They showed that this ensemble of FI techniques, coupled with multiple ML

approaches, produced more robust importance compared to single FI methods. This framework has been used to identify features relevant to the creep rate of additive manufactured materials [25].

2.3. Limitations of existing methods

These single or multiple FI techniques suffer from several limitations. They have limited interpretability; the methods and outputs are difficult for domain experts to understand. For example, what does a final FI coefficient of 0.76 mean (is it a good high value, or a low value)? The value obtained also depends on the FI method adopted and those that are not model agnostic also depend on the specific ML model(s) employed. Use of crisp fusion can also result in significant loss of information as data subspaces and gradients are not adequately explored because of the way training/test sampling is conducted. Additionally, crisp ensembles like majority voting discard extreme values of FI which, for safety critical systems, should not be disregarded. Finally, there is a paucity of literature that explores data sampling for FI. Using solely training or test sets for FI leads to loss of importance information, as fewer combinations of data instances are investigated. Here, we show how these limitations can be alleviated using a fuzzy logic approach for representation of importance. This accounts for multiple levels of uncertainty, whilst simultaneously simplifying interpretation of results. Fuzzy logic defines data/context-dependent intervals of importance that make it easier to identify data regions where specified features have high or low importance for decision making. This level of interpretation granularity will accelerate data-driven intelligent research and decisions by identifying critical data instances that enhance prognosis and health management, manufacturing, medicine and design of new materials.

3. Methods

3.1. Fuzzy ensemble feature importance

Our approach uses an ensemble of ML models coupled with multiple FI techniques to generate a large dataset of FI values. This dataset is analysed by a fuzzy logic (FL) system that specifies, for each ML and for each feature, low, moderate, or high importance. A flowchart of our framework is shown in Fig. 1 and the four stages in the framework are described in detail in the following sections.

3.1.1. Data pre-processing

The data is pre-processed to manage missing values and outliers, identify relevant features, and to perform data normalisation. Feature reduction, engineering, or selection can be used to remove redundant, low variance, and less relevant features, reducing noise in the data. The pre-processed data are subsequently transformed to a format suitable for training the ML models.

The pre-processed data is partitioned into k equally sized subsets for model training using k -fold cross validation [26] or any other type of bootstrapping. In each training step, 1 subset is held out and the remaining $k-1$ subsets are used to train the model. The training process is repeated k times, holding out a different subset each time. The k models that result allows the variance in feature gradients and feature contributions to be calculated.

3.1.2. Machine learning predictions

An ensemble of models is generated using n different ML algorithms, each algorithm being trained k times on the partitioned subsets of data. It is important to note that the framework can be extended to include any type of ML method and model-agnostic feature importance technique. This produces $k * n$ trained

models. We optimise hyperparameters to ensure that the most accurate models are obtained for the specified problem.

3.1.3. Feature importance

Each trained model employs a subset of j FI techniques and a validation data subset to produce a large set of FI coefficients for each feature and importance method (see Fig. 2). The dataset of FI coefficients encompasses the variance in the learning characteristics and performance of ML approaches, the variance in calculating importance by the various FI techniques, and the variance of feature contributions in the data.

3.1.4. Fuzzy interpretation

Here, we generate MFs and inference rules from the data obtained from stage 3. These elucidate whether a feature is of low, moderate, or high importance for each ML method used and how important a feature is after fusing importances from the different approaches.

(a) Generating membership functions: The data set of FI coefficients is normalised to the same scale to ensure unbiased and consistent representation of importance. ML MFs are generated using the FI coefficients from each ML method (i.e. each 'Model' column in Fig. 2 will produce a MF), representing the range of importance in the feature set. Feature MFs are generated by combining the ML MFs for each feature (each 'Feature' row and each 'Model' column will produce MFs) using the minimum (min) and maximum (max) operators. The min operator combines the lower bounds of the MFs and the max operator combines the upper bounds. The feature MFs represent how important a feature is relative to other features for each ML method. The output MF (final importance MF) is generated by mapping the FI terms obtained from the ML MFs of the predicted FI values to the ML MFs of the actual (ground truth) FI coefficients for the entire feature set. The output MF presents the final importance of features after combining their importance from the different ML approaches. The MFs (ML, feature, and output MFs) consist of the linguistic fuzzy sets (FSs) 'low', 'moderate' or 'high' importance obtained from boxplot distributions of the FI coefficients, as proposed by Mase et al. [5]. The five class summary (minimum (min), first quartile (Q_1), median (Q_2), third quartile (Q_3) and maximum (max)) from the boxplot distribution are used to construct the MFs. We use Z and S-MFs for 'low' and 'high' FSs as they are most suitable for representing the extreme boundaries with maximum degrees or likelihood of membership. For instance, we expect a feature's importance coefficient of 0 to be in the 'low' FS with a degree of 1 (maximum likelihood) and an importance coefficient of 1 to be in the 'high' FS with a degree of 1. The Z-MF is generated using the minimum and median values of the distribution and the S-MF employs the median and maximum values. For 'moderate' FS, we use Triangular MFs as the most efficient way to represent intermediate FSs for our problem. We use first quartile, median and third quartile values to construct triangular-MFs.

Fig. 3 shows an example of MFs generated from a boxplot distribution. Fig. 3(a) represents the distribution of importance coefficients for the Gradient Boost approach after training several models on different data samples and applying suitable FI techniques on the trained models as described in Stages 2 and 3 (Fig. 1). Using the five-summary values of the boxplot (min = 0, $Q_1=0$, $Q_2=0.18$, $Q_3=0.7$ and max = 1), we obtain the MFs in Fig. 3 (b). 'low' MF = Z(0, 0.18), 'moderate' MF = triangular (0, 0.18, 0.7) and 'high' MF = S(0.18, 1). The MFs provide a clear representation of the range and variation of importance coefficients in the low, moderate and high categories. For example, the low category in Fig. 3(b) consists of importance coefficients from 0 to 0.18, a lower variation of values compared to moderate and high categories as shown by its narrow support (range). (b) Fuzzy rule generation: After developing the MFs, we employ the Wang-Mendel

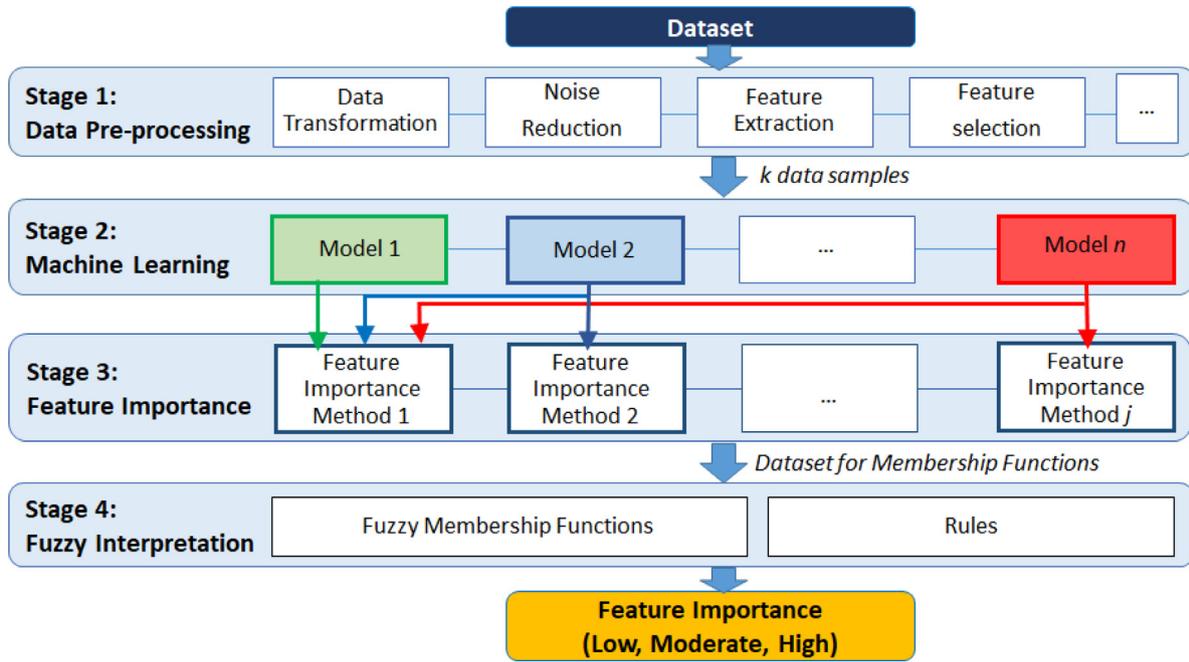


Fig. 1. The fuzzy feature importance framework.

Data for Membership Functions of Importance per Machine Learning Model

		Model 1		Model 2		Model n	
		Method 1 Importance		Method 1 Importance	Method 2 Importance	Method 1 Importance	Method j Importance
Data For Membership Functions of Individual Features' Importance	Sample 1	FI Feature 1	FI Feature 1	FI Feature 1	FI Feature 1	FI Feature 1	FI Feature 1
	Feature 1
	Sample k	FI Feature 1	FI Feature 1	FI Feature 1	FI Feature 1	FI Feature 1	FI Feature 1

	Feature i

Range of importance of a feature for the problem investigated

Range of Importance found within the feature set for a Machine Learning approach

Fig. 2. Data from Stage 3.

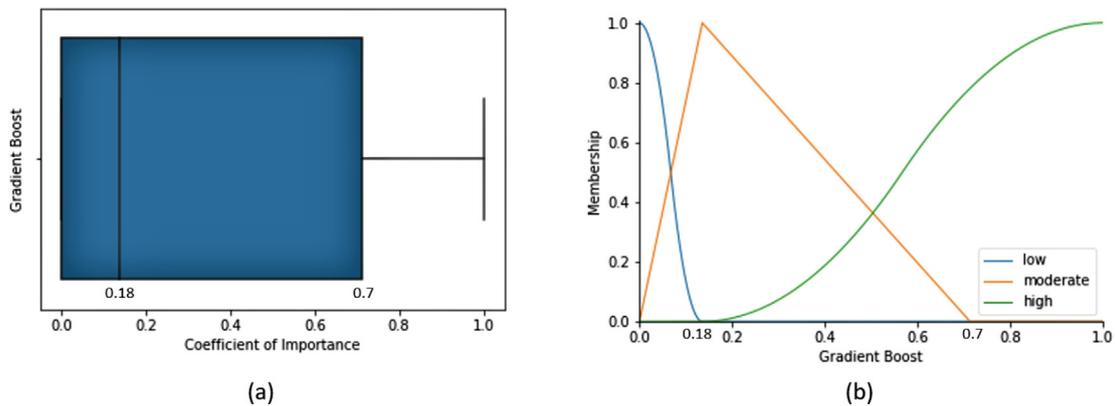


Fig. 3. Membership functions generated from boxplot distribution: a) boxplot distribution of importance values, and b) extracted MFs using our boxplot data-driven approach.

[17] approach to learn rules that map the FI coefficients from the different ML approaches (inputs) to the final description of importance (output). The Wang-Mendel is most applicable for ground truth data as it uses a supervised learning approach to capture

the appropriate mappings between the inputs and output. However, it is important to present other available unsupervised and supervised rule generation methods, such as fuzzy C-means Wang-Mendel [27], mountain clustering [28], fuzzy decision trees

[29] and genetic algorithms [30]. Using the data structure in Fig. 2, each feature i has a ground truth importance coefficient, and each feature sample k has a set of FI coefficients obtained by applying multiple (ML, FI) pairs. For example, feature 1 - sample 1 in Fig. 2 will produce the following input–output pairs:

$$\begin{aligned} &(F_{ML_1}^{1,1}, F_{ML_2}^{1,1}, F_{ML_n}^{1,1}; y^1), \\ &(F_{ML_2}^{1,2}, F_{ML_n}^{1,2}; y^1), \\ &\vdots \\ &(F_{ML_n}^{1,j}; y^1), \end{aligned} \tag{1}$$

where $(F_{ML_i}^{1,j})$ denotes the corresponding importance coefficient of Feature 1 obtained from ML model 1 and FI method 1 and y^1 denotes the ground truth importance coefficient for Feature 1. Each input–output pair is assigned to the corresponding ML and output MFs to produce the rules and the number of occurrences of each rule (weight). The rules show mappings of FI terms obtained from the different ML models to the final FI description, while the rule weights are utilised to implement a rule reduction procedure for conflicting rules. For example, Table 1 shows sample rules to demonstrate how conflicting rules could be resolved using rule weights. We observe that rule 2 and rule 5 are in conflict, with weights 5 and 2 respectively. We eliminate rule 5 as it has a smaller weight. Other rule reduction procedures, such as evidence combinational models [31] may be applied in this stage. The rules could be further improved by domain experts due to their simple linguistic representation. For a more elaborate description of the Wang-Mendel rule generation method, please refer [17,31].

(c) Fuzzy inference system: The rules utilised during inference fuse the importance of a feature from multiple ML approaches to obtain a final importance (low, moderate, high). To achieve this, we employ the Mamdani inference technique [18], a rule-based fuzzy inference method. The technique converts FI coefficients into FSs (fuzzification) using the MFs generated in Stage 4 (a), computes rule strengths, and uses these to estimate feature importance as low, moderate, or high. The rule strengths are computed using ‘AND’ (minimum of membership degrees), ‘OR’ (maximum of membership degrees) and ‘NOT’ (1 minus membership degree) [32]. We use the following sample rule to illustrate the calculation of rule strengths and final FI importance:

Rule 1: IF ML_1 is ‘low’ AND ML_2 is ‘low’ AND ML_n is ‘low’ THEN Output is ‘low’

Assume ML_1, ML_2, ML_n have the MFs found in Fig. 4 and the FI coefficients produced for feature 1 by ML_1 is 0.1, ML_2 is 0.2, and ML_n is 0.1 using FI Method 1. The fuzzification process produces the following degrees of membership for the FSs in Rule 1 (as shown in Fig. 5):

$$\begin{aligned} ML_1(\text{low}) \text{ for } 0.1 &= 0.1 \\ ML_2(\text{low}) \text{ for } 0.2 &= 0.7 \\ ML_n(\text{low}) \text{ for } 0.1 &= 0.7 \end{aligned}$$

Using the rule operators in Rule 1 (AND), we obtain the rule strength = $ML_1(\text{low})$ AND $ML_2(\text{low})$ AND ML_n = $\min[0.1, 0.7, 0.7]$ = 0.1.

The final step in the Mamdani inference method is to aggregate rule strengths of the entire rule set produced by the rule generation method. This is achieved by finding the maximum rule strength [18] for rules with similar output FSs (e.g. the ‘THEN’ part of rules). The aggregated rule strengths represent the likelihood of the final FI belonging a specific category. For example, if the data from Stage 3 of the framework produce six rules with the following rule strengths and output FS for the FI coefficients of feature 1:

Table 1
An illustration of conflicting rules with different rule weights (i.e. Rules 2 and 5).

Rule No	Rule Description	Weight
1	IF ML_1 is ‘low’ and ML_2 is ‘low’ and ML_n is ‘low’ THEN Output is ‘low’	5
2	IF ML_1 is ‘high’ and ML_2 is ‘high’ and ML_n is ‘high’ THEN Output is ‘high’	5
3	IF ML_1 is ‘moderate’ and ML_2 is ‘moderate’ and is ML_n ‘low’ THEN Output is ‘low’	4
4	IF ML_1 is ‘moderate’ and ML_2 is ‘moderate’ and ML_n is ‘high’ THEN Output is ‘high’	4
5	IF ML_1 is ‘high’ and ML_2 is ‘high’ and ML_n is ‘high’ THEN Output is ‘moderate’	2

- Rule1 : Strength = 0.1, Output = low
- Rule2 : Strength = 0, Output = low
- Rule3 : Strength = 0.2, Output = moderate
- Rule4 : Strength = 0.1, Output = moderate
- Rule5 : Strength = 0.5, Output = high
- Rule6 : Strength = 0.8, Output = high

The final FI of feature 1 using the above rules is:

- $\max(0.1,0.0)=10\%$ likelihood of low importance
- $\max(0.2,0.1)=20\%$ likelihood of moderate importance
- $\max(0.5,0.6)=60\%$ likelihood of high importance

3.2. Data generation

We validate our approach using synthetic data described by Rengasamy et al. [4], as it allows us to verify the efficacy of feature importance produced by FEFI against the ground truth. Nine datasets are used for regression, with 2000 instances and different number of features, feature informative levels, feature noise levels, or feature interaction levels. Feature informative levels affects the percentage of features contributing to the output data. For example, for ten features with 50% feature informative levels indicates that five features affects the data output while the remaining five features is independent of the output. Subsequently, the noise levels in features are adjusted using the Gaussian noise standard deviation. Larger values push feature values further from the average, decreasing Signal/Noise ratios. Feature interaction strengths are an additional characteristic not investigated by Rengasamy et al. Feature interaction levels of data determines the strength of interaction between features. It can be varied by changing the effective rank of the matrix, i.e., the maximum number of linearly independent features in a matrix. If a matrix has three columns (features) and the first two columns are linearly independent, the matrix has a rank of at least 2. The feature matrix ranks are generated by the `make_regression` function in Python `scikit-learn` library [33]. Fig. 6 shows the correlation matrix for the three datasets with fully independent features compared to those with highly correlated features. The correlation between features is calculated using Pearson correlation.

Additionally, different feature interaction intensities are added to each dataset, as illustrated in Fig. 7. Feature interaction occurs when a feature affects one or more features, making it difficult to discern its actual FI. The interaction effect between pairs of features is calculated using Friedman’s H statistics [34]. Table 3 shows summary of different datasets tested in this paper. The dataset’s properties are altered one at a time.

We use partial dependence plots (PDPs) to analyse and illustrate the relationship between the features and output. It shows that for high feature interaction strengths the output response sur-

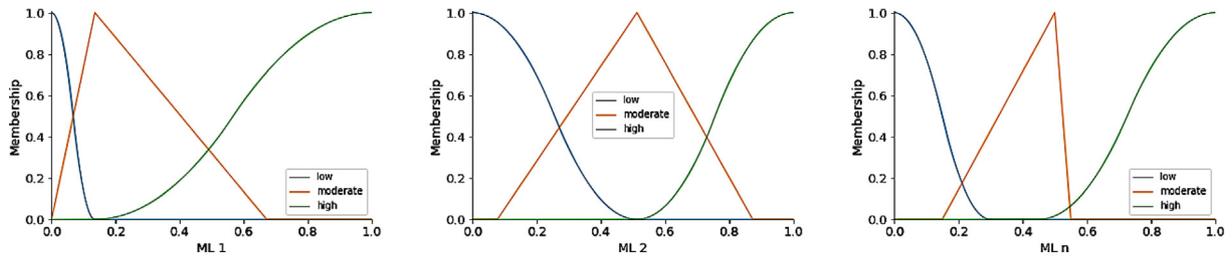


Fig. 4. Sample membership functions for FI coefficients produced by ML approaches.

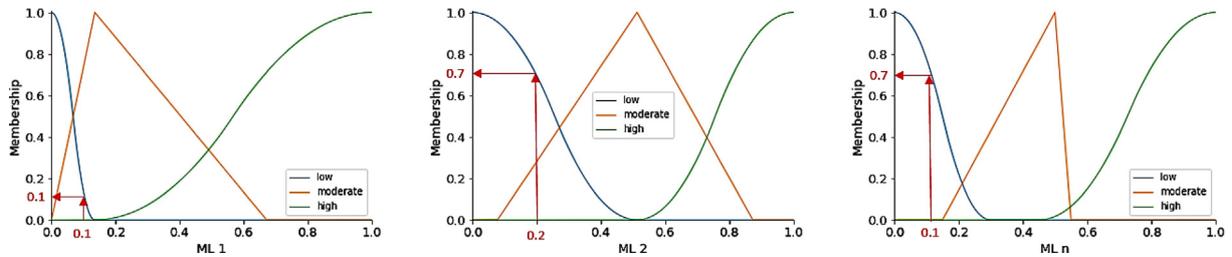


Fig. 5. Fuzzification of FI coefficients using Rule 1 and ML membership functions in Fig. 4.

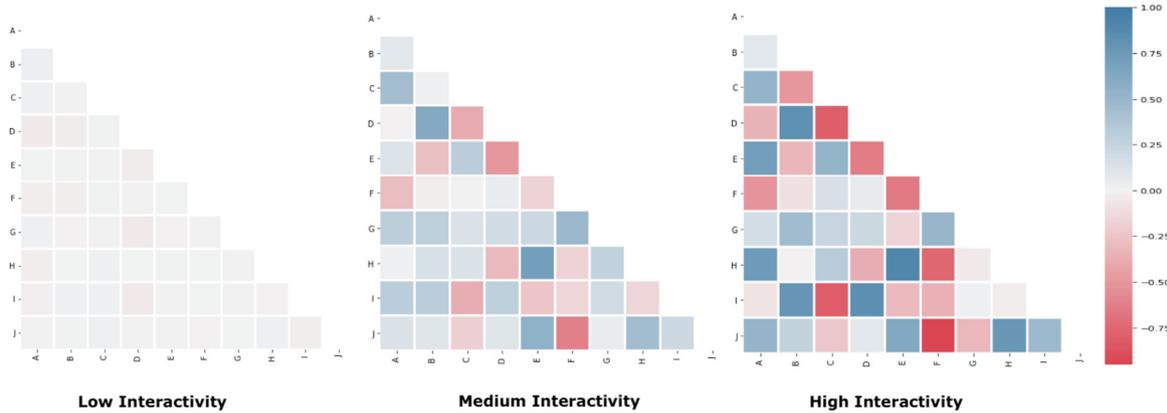


Fig. 6. Pearson correlation matrix for datasets with different interaction strength between ten synthetic features investigated, labelled from A to J.

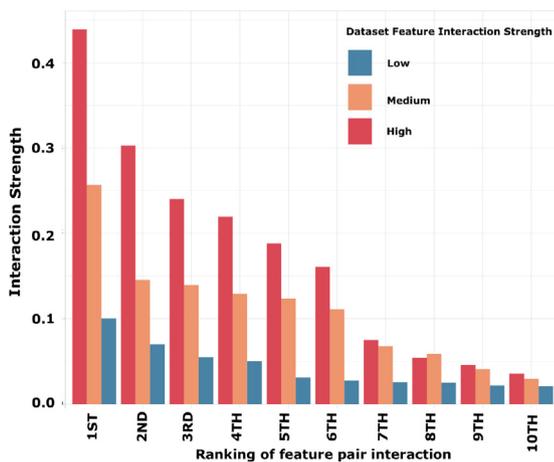


Fig. 7. Interaction strength between pairs of features for each of the three datasets. The x-axis denotes the feature pairs with the top 10 highest interaction from each of the datasets.

face may not be linear. The PDPs show the marginal influence of one or two combined features on the output using Monte Carlo

Table 3
Datasets with different feature interaction intensities, number of features, features informative level, and noise's standard deviation tested in the FEFI framework.

Dataset properties	Values of properties
Number of features	10, 30, 50
Feature interaction	Low, Medium, High
Features informative level	20%, 50%, 90%
Noise's standard deviation	0.5, 2.0, 5.0

method across the dataset [20]. The PDP in Fig. 8 illustrates the non-linear relationship between two features and the output in a high interactive region of the dataset.

3.3. Machine learning methods selection and configurations

Five ML methods were implemented: random forest (RF), gradient boost (GB), extra trees (ET), support vector regression (SVR), multiLayer perceptron (MLP). Three FI methods: PI, SHAP, and Gini Importance are used. Diverse ML and FI techniques are chosen to show the benefit of FEFI in modelling the uncertainties in FI due

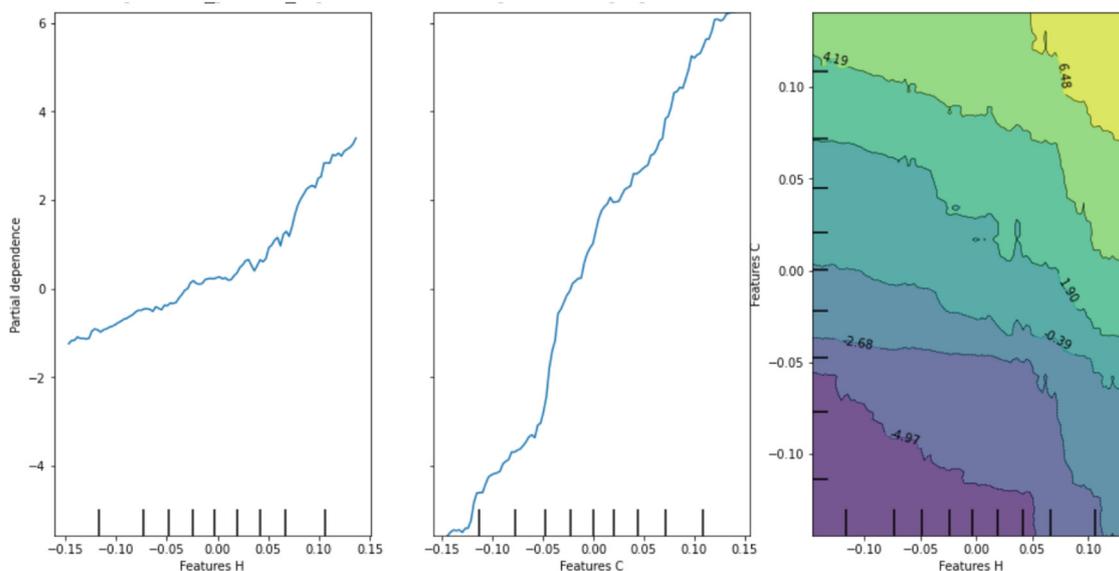


Fig. 8. The partial dependence plot of two features on the output. The left and the center plot shows the respective partial dependence of the respective feature to output. The right plot shows the combined response on the output for the two features.

to differences in learning and FI identification characteristics. The hyperparameters for each ML method are shown in Table 4.

3.4. Evaluation metrics

The FEFI approach is evaluated and compared to mean and majority vote EFI using 4 different data properties: features interaction, number of features, features informative level, and noise level of features. The mean absolute error (MAE) and root mean squared error (RMSE) metrics compared these properties to the ground truth feature importance. Three of the four data properties, except feature interaction were tested by Rengasamy et al. [4]. Each data property was further split into 3 different subsets: training; test; and whole (train + test) set for feature importance quantification and comparisons between mean EFI, majority vote EFI, and FEFI.

4. Results and discussion

4.1. Results

Here we compare the results of FEFI to the best crisp fusion strategies from Rengasamy et al. [4]; the mean and majority vote across the different datasets shown in the Experimental design section (Table 2). Majority vote crisp strategy first ranks features based on their importance coefficients produced by each ML model and FI technique. Later, it averages the coefficients in the most common ranks for each feature. For example, if a feature has the set of ranks (1,2,1,3) amongst 10 features for coefficients produced by 4 ML models, the final importance of the feature will be the average of the coefficients with the rank equals to 1 as it is the most common rank in the set of ranks. The mean crisp strategy simply averages the importance coefficients for each feature produced by the ML models and FI techniques.

Table 5 shows the errors in calculating the final FI using the different fusion strategies as the number of features in the data increases (datasets 1, 2 and 3 in Table 2). Increasing the number of data features with a 90% feature informative level made it more difficult for crisp strategies to quantify feature importance. This is illustrated by the significant increase in FI errors from 10 to 30 features [p-value < 0.001 for all data subsets] and from 10 to 50 fea-

Table 4 Hyperparameters for each of the machine learning approaches employed to test the Fuzzy Ensemble Feature Importance (FEFI) framework.

ML Approaches	Hyperparameters	Values
Gradient Boosting Regressor	Loss	Least squares
	Learning rate	0.1
	Number of boosting stages	50
	Splitting criterion	Friedman MSE
	Minimum samples to split	2
Random Forest Regressor	Maximum depth	3
	Number of Trees	50
	Max depth	None
	Splitting Criterion	MSE
	Minimum samples to split	2
Extra Trees Regressor	Bootstrap	True
	Number of Trees	50
	Max depth	None
	Splitting Criterion	MSE
	Minimum samples to split	2
Support Vector Regressor	Bootstrap	False
	Kernel	Linear
	Tolerance	0.001
	Regularisation	1.0
	Epsilon	0.1
MultiLayer Perceptron	L2 Regulariser	0.0001
	Learning Rate	0.001
	Hidden Layer Size	50
	Activation function	ReLU
	Optimiser	Adam

Table 2 Datasets with different feature interaction, number of features, features informative level, and noise's standard deviation (std) tested on the proposed Fuzzy Ensemble Feature Importance methods.

Dataset Number	Number of features	Feature interaction	Feature informative level	Feature noise std
1	10	Low	90%	0.5
2	30	Low	90%	0.5
3	50	Low	90%	0.5
4	10	Moderate	90%	0.5
5	10	High	90%	0.5
6	10	Low	20%	0.5
7	10	Low	50%	0.5
8	10	Low	90%	2.0
9	10	Low	90%	5.0

Table 5

MAE and RMSE comparison between three different frameworks for feature importance fusion/ensemble: (1) Mean, (2) Majority Vote, (3) FEFI on three datasets with different number of features using three data subsets (best performance in bold).

Number of features	Data subset	Mean (crisp)		Majority Vote (crisp)		FEFI	
		MAE	RMSE	MAE	RMSE	MAE	RMSE
Dataset 1: 10	Whole	0.146	0.168	0.124	0.143	0.148	0.181
	Train	0.146	0.168	0.124	0.143	0.126	0.156
	Test	0.154	0.181	0.128	0.153	0.135	0.169
Dataset 2: 30	Whole	0.379	0.435	0.397	0.454	0.107	0.133
	Train	0.379	0.434	0.397	0.454	0.109	0.135
	Test	0.377	0.433	0.398	0.455	0.117	0.148
Dataset 3: 50	Whole	0.297	0.365	0.293	0.360	0.128	0.155
	Train	0.298	0.365	0.294	0.361	0.131	0.155
	Test	0.286	0.354	0.284	0.351	0.128	0.155

tures [p-value < 0.001 for all data subsets]. This is due to the inefficiency of crisp methods in capturing the increased variation of FI coefficients produced by additional features. In contrast, FEFI exhibited no significant differences in FI errors for different number of features [p-value > 0.1 for all data subsets] and significantly lower FI errors with 30 and 50 data features compared to the crisp mean and majority vote [p-value < 0.001 for all data subsets]. This indicates that FEFI efficiently captures the increased FI uncertainty caused by extra number of features compared to the crisp strategies.

Table 6 shows errors in calculating final FI as the interactions between features in the data increase (datasets 1, 4 and 5 in Table 2). As with the number of features, there is a significant increase in errors for mean and majority vote fusion strategies when the interactions between data features increase from low to moderate [p-value<0.05 for all data subsets] and from low to high interaction [p-value<0.02 for all data subsets]. However,

there is no significant increase in errors from medium to high feature interaction [p>0.05 for all data subsets]. For FEFI, we again observe no significant difference in FI errors for different levels of feature interaction [p>0.1 for all data subsets] and significantly lowers FI errors for medium and high feature interaction levels compared to mean and majority vote [p<0.001 for all data subsets]. Again, this indicates that the fuzzy approach efficiently captures the increased variation of FI coefficients caused by higher feature interactions compared to the crisp strategies.

Table 7 presents the FI errors obtained from the different fusion strategies as the informative level of features increases i.e. datasets 1, 6 and 7 in Table 2. We observe a significant rise in FI errors for mean and majority vote ensembles as feature informative levels increase from 20% to 50% [p<0.001 for all data subsets], 20% to 90% [p<0.0001 for all data subsets] and 50% to 90% [p<0.005 for all data subsets]. This means the variation between the calculated FI coefficients increases as the informative level of features

Table 6

MAE and RMSE comparison between three different frameworks for feature importance fusion/ensemble: (1) Mean, (2) Majority Vote, (3) FEFI on three different features interaction level using three data subsets (best performance in bold).

Interaction level	Data subset	Mean (crisp)		Majority Vote (crisp)		FEFI	
		MAE	RMSE	MAE	RMSE	MAE	RMSE
Dataset 1: Low	Whole	0.146	0.168	0.124	0.143	0.148	0.181
	Train	0.146	0.168	0.124	0.143	0.126	0.156
	Test	0.154	0.181	0.128	0.153	0.135	0.169
Dataset 4: Medium	Whole	0.203	0.248	0.220	0.275	0.141	0.189
	Train	0.203	0.249	0.221	0.276	0.135	0.179
	Test	0.202	0.251	0.216	0.273	0.148	0.188
Dataset 5: High	Whole	0.222	0.288	0.249	0.314	0.125	0.164
	Train	0.222	0.288	0.250	0.315	0.139	0.177
	Test	0.223	0.291	0.226	0.288	0.144	0.180

Table 7

MAE and RMSE comparison between three different frameworks for feature importance fusion/ensemble: (1) Mean, (2) Majority Vote, (3) FEFI on three different features informative level using three different subsets of data (best performance in bold).

Informative level	Data subset	Mean (crisp)		Majority Vote (crisp)		FEFI	
		MAE	RMSE	MAE	RMSE	MAE	RMSE
Dataset 6: 20%	Whole	0.008	0.011	0.007	0.009	0.078	0.177
	Train	0.008	0.011	0.007	0.009	0.084	0.178
	Test	0.010	0.012	0.008	0.010	0.078	0.178
Dataset 7: 50%	Whole	0.078	0.101	0.061	0.086	0.111	0.164
	Train	0.079	0.102	0.062	0.088	0.099	0.144
	Test	0.093	0.134	0.078	0.121	0.111	0.164
Dataset 1: 90%	Whole	0.146	0.168	0.124	0.143	0.148	0.181
	Train	0.146	0.168	0.124	0.143	0.126	0.156
	Test	0.154	0.181	0.128	0.153	0.135	0.169

increases. However, this increase in the variation of FI coefficients caused by an increase in feature informative levels is not significant compared to the increase caused by feature numbers, noise and feature interaction levels. This can be observed in the inferior performance of FEFI for all three informative levels compared to the crisp fusion methods. Therefore, crisp fusion strategies that produce clear-cut decision boundaries become superior to fuzzy approaches when the importance of features is distinct. Nonetheless, the results for FEFI show no significant increase in FI errors as feature informative levels increase from 20% to 50% [$p > 0.1$ for all data subsets], 20% to 90% [$p > 0.01$ for all data subsets] and 50% to 90% [$p > 0.05$ for all data subsets]. This implies FEFI shows better performance in capturing increased FI variation produced by an increase in feature informative level.

Table 8 illustrates the errors obtained from calculating the final FI as the level of noise in data increases from 0.5 noise standard deviation to 2.0 and 5.0 (datasets 1, 8 and 9 in Table 2). We observe no significant difference in FI errors for mean, majority vote and FEFI fusion strategies as the noise standard deviation increases

from 0.5 to 2.0 [$p\text{-value} > 0.5$ for all data subsets], 0.5 to 5.0 [$p > 0.5$ for all data subsets] and 2.0 to 5.0 [$p\text{-value} > 0.9$ for all data subsets]. This means the variation or uncertainty of FI coefficients is not affected by changes of the level of noise in the data. However, FEFI produces a significantly lower FI error compared to the crisp mean and majority vote when the noise standard deviation is 5.0 [$p < 0.05$ for whole and test subsets]. This is because of its capability to effectively capture noise in data using its MFs.

In summary, the results show that our fuzzy FI fusion approach outperforms mean and majority vote feature importance fusion methods in capturing increased variation of FI coefficients caused by increased data dimensionality, complexity and noise. This is because FEFI explores the data space more thoroughly as it uses multiple samples of data to make decisions about the importance of features. It also uses distributions of the data to provide better definitions of FI and soft boundaries to capture the intermediate uncertainties of FI classification.

Finally, it is important to mention that majority vote generates apparently lower errors compared to FEFI for 10 features, low fea-

Table 8

MAE and RMSE comparison between three different frameworks for feature importance fusion/ensemble: (1) Mean, (2) Majority Vote, (3) FEFI on three different features noise levels using three different subsets of data (best performance in bold).

Noise's Standard Deviation	Data subset	Mean (crisp)		Majority Vote (crisp)		FEFI	
		MAE	RMSE	MAE	RMSE	MAE	RMSE
Dataset 1: 0.5	Whole	0.146	0.168	0.124	0.143	0.148	0.181
	Train	0.146	0.168	0.124	0.143	0.126	0.156
	Test	0.154	0.181	0.128	0.153	0.135	0.169
Dataset 8: 2.0	Whole	0.150	0.170	0.150	0.169	0.125	0.152
	Train	0.150	0.171	0.152	0.169	0.133	0.180
	Test	0.165	0.195	0.163	0.200	0.131	0.161
Dataset 9: 5.0	Whole	0.151	0.173	0.155	0.180	0.117	0.137
	Train	0.150	0.172	0.151	0.171	0.131	0.168
	Test	0.164	0.195	0.165	0.197	0.117	0.137

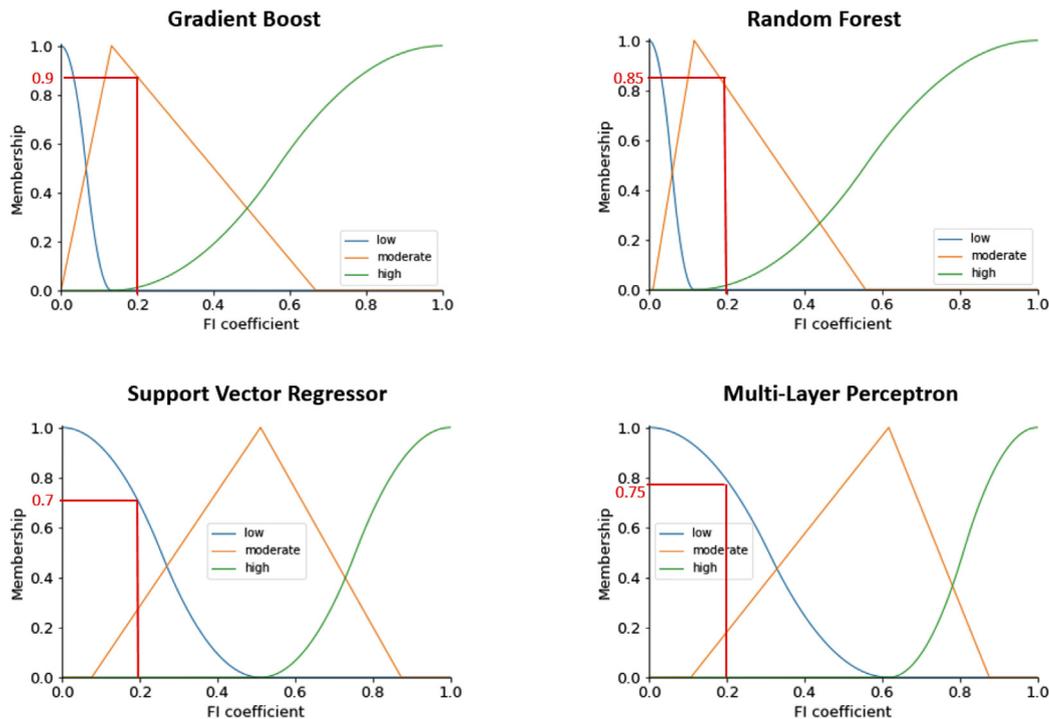


Fig. 9. Membership functions of ML approaches generated from Dataset 1 showing different interpretations of importance for the same FI coefficient.

ture interaction level, 0.5 noise standard deviation, 20%, 50% and 90% informative levels. However, the difference in errors is not statistically significant [$p > 0.05$ for all data subsets] except for 20% and 50% feature informative levels where $p < 0.001$.

4.2. Discussion

We now discuss the interpretability, advantages and limitations of FEFI in determining importance of features. FEFI provides clear and accurate interpretation of feature contribution to different ML model outcomes using the ML MFs. For example, Fig. 9 shows different interpretations for the same FI coefficient (0.2) by the MFs generated from dataset 1 (refer dataset description in Table 2) for GB, RF, SVR, and MLP. We observe that a FI coefficient of 0.2 has a 70% likelihood of low importance for SVR and a 75% likelihood of low importance for MLP, similar to what we would expect for a coefficient close to 0. However, for GB and RF, a coefficient of 0.2 has a high likelihood of moderate importance and no likelihood of low importance i.e. 90% likelihood of moderate importance for GB and 85% for RF. Therefore, without the generation of these MFs, the interpretation of the coefficients produced by the different ML approaches will be misleading. In addition, Fig. 9 supports our use of an ensemble of diverse ML approaches to complement outputs. We observe different MF representations for the different ML approaches but similar MFs for approaches with similar learning characteristics. For example, the low importance FS for GB is very narrow compared SVR and MLP but very similar to that of RF due to similar tree-based learning process.

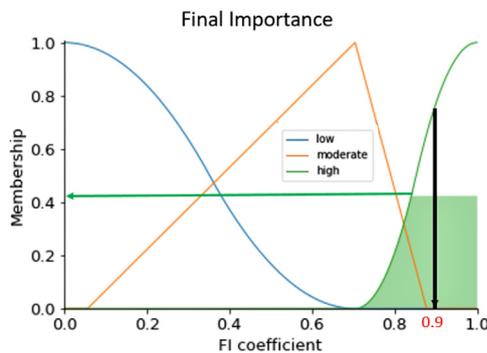


Fig. 10. An example of the final importance of a feature after fusing the importance coefficients obtained from the different ML approaches.

FEFI also provides more reliable interpretations of the final FI for domain experts. Using the output MFs as crisp FI coefficients produced by fusion strategies may be misleading without context. For example, Fig. 10 shows the final importance coefficient (0.9) after fusing the importance obtained from the different ML approaches. Without context (the output MF), it might be assumed that 0.9 has a high likelihood of high importance as 0.9 is very close to 1 (maximum likelihood), whereas, we find that 0.9 has a 40% likelihood of high importance when the entire feature space is considered (indicated by the height of the shaded region in Fig. 10). This context provided by FEFI is important in safety-critical and medical diagnosis in providing the extent or critical level of a feature’s contribution to the decision.

Furthermore, FEFI shows great potential as a diagnostic tool in identifying data subsets with extreme cases of feature importance (this might be the case of activity cliffs in quantitative structure activity relationship modelling) or with significant variation of feature importance (e.g., for significantly nonlinear response surfaces [35]). This is achieved by producing feature MFs at different data samples or time steps to investigate the importance of a feature in the samples. For example, Fig. 11 shows the importance of a feature generated by different data samples of dataset 1 when PI is used to calculate the feature’s importance for best performing GB model. We observe that the feature’s importance coefficient is different in the samples; 0.8 in sample 20 with a 90% likelihood of high importance and 0.5 in sample 60 with a 30% likelihood of high importance. Real-world systems can have more extreme cases as they produce dynamic and multifaceted data where the importance of a feature may vary in different samples due to interactions with other features or supplementary data from other data sources. For example, in healthcare, when multiple sources of data (patient surveys, clinical trials etc) associated with a heterogeneous patient population are merged, features may have different importance depending on patient stratification.

In real-world systems where the ground truth FI is unavailable, complete automation is limited using FEFI. This is because development of the user-friendly IF/THEN rules requires: interaction with ML experts based on their experience using ML approaches and their expectations of the ML performance on the data; a comprehensive literature review to find out the rules describing the generation of FI by the different ML approaches; or the use of heuristic rule generation methods such as genetic algorithms [36]. In addition, the number of rules increases exponentially with an increase in the number of ML approaches, making it difficult for experts to provide rules and reduces the interpretability of the rules.

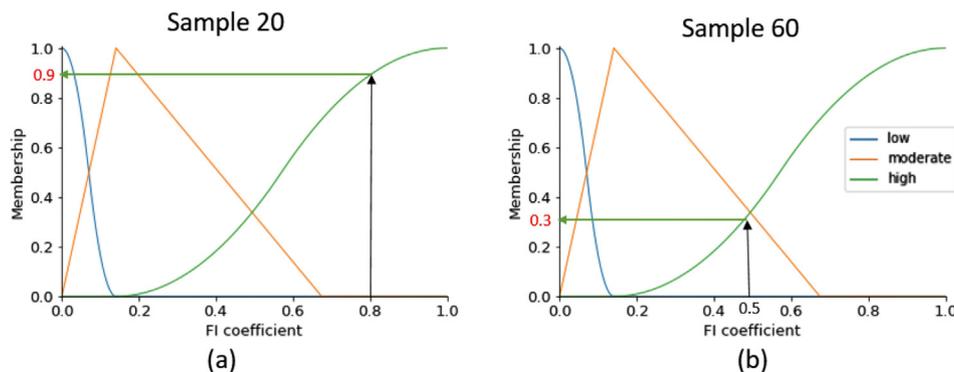


Fig. 11. The same feature having different FI coefficients in different samples of Dataset 1. (a) Sample 20 produces a FI coefficient of 0.8, and (b) Sample 60 produces a FI coefficient of 0.5.

5. Conclusions and future work

This paper introduces a FEFI framework to interpret ensemble of ML methods and FI methods. This is primarily motivated by the uncertainty and unreliable of post-processing FI methods along with the ethical implication of untrustworthy output interpretations. To overcome these issues we proposed a framework that combines ensemble feature importance methods with fuzzy logic to utilise the distributions of the FI output to provide better definitions of FI and soft boundaries to capture the intermediate uncertainties of FI classification. Furthermore, Fuzzy Systems have the added advantage of explaining the FI in linguistic term making it easier for non-experts to understand the implication of FI outputs. The combination of robust uncertainty estimations and simple linguistic representation of the FI outputs allows users to make informed decisions – this is especially important for safety critical applications. We applied the methodology to 9 different synthetic datasets all with varying level of number of features, feature interaction strengths, informative level, and noise level. The usage of synthetic datasets allow us to empirically measure the accuracy of FEFI framework. The results shows that by considering the uncertainty produced in the multi-ML and multi-FI methods settings, FEFI was able to produce more accurate FI outputs compared to crisp mean and majority vote-based ensemble FI frameworks.

For future work, we intend to evaluate our methodology using other ML methods such as neural networks with more complex learning processes and compare our Fuzzy rule-based approach with Fuzzy integral-based approach, which is also a promising fuzzy fusion strategy. Additionally, there are multiple ways of capturing uncertainty in data besides fuzzy logic, such as Bayesian statistics. A comparison of the framework using fuzzy logic and Bayesian statistics would provide a greater insight into which uncertainty estimation method is more suitable for providing a more robust and accurate FI interpretation for the end users.

Lastly, we plan to perform further investigations and evaluations of our approach on real-world datasets. This is important in revealing insights about the performance of our approach on real-world scenarios.

CRedit authorship contribution statement

Divish Rengasamy: Conceptualization, Methodology, Software, Formal analysis, Data curation, Writing - original draft, Writing - review & editing, Project administration. **Jimiamma M. Mase:** Conceptualization, Methodology, Software, Formal analysis, Data curation, Writing - original draft, Writing - review & editing, Project administration. **Aayush Kumar:** Software. **Benjamin Rothwell:** Conceptualization, Supervision. **Mercedes Torres Torres:** Conceptualization, Methodology, Validation, Writing - original draft. **Morgan R. Alexander:** Writing - original draft, Writing - review & editing. **David A. Winkler:** Methodology, Validation, Writing - original draft, Writing - review & editing. **Grazziela P. Figueredo:** Conceptualization, Methodology, Validation, Writing - original draft, Writing - review & editing, Supervision.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

The second author is supported by the Horizon Centre for Doctoral Training at the University of Nottingham (UKRI Grant No. EP/L015463/1) and by Microlise.

References

- [1] Sandeep Reddy, Sonia Allan, Simon Coghlan, Paul Cooper, A governance model for the application of AI in health care, *J. Am. Med. Inform. Assoc.: JAMIA* 27 (3) (2019) 491–497.
- [2] Felix Gille, Anna Jobin, Marcello Ienca, What we talk about when we talk about trust: Theory of trust for AI in healthcare, *Intell.-Based Med.* 1–2 (2020) 100001.
- [3] Divish Rengasamy, Mina Jafari, Benjamin Rothwell, Xin Chen, Graziela P Figueredo, Deep learning with dynamically weighted loss function for sensor-based prognostics and health management, *Sensors* 20 (3) (2020) 723.
- [4] Divish Rengasamy, Benjamin C. Rothwell, Graziela P, Figueredo, Towards a more reliable interpretation of machine learning outputs for safety-critical systems using feature importance fusion. *Appl. Sci.*, 11(24):11854, 2021.
- [5] Jimiamma Mafeni Mase, Utkarsh Agrawal, Direnc Pekaslan, Mohammad Mesgarpour, Peter Chapman, Mercedes Torres Torres, and Graziela P Figueredo. Capturing uncertainty in heavy goods vehicles driving behaviour. In 2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC), pages 1–7. IEEE, 2020.
- [6] Alejandro Barredo Arrieta, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Bénéto, Siham Tabik, Alberto Barbado, Salvador García, Sergio Gil-López, Daniel Molina, Richard Benjamins, et al. Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai. *Information Fusion*, 58:82–115, 2020.
- [7] Utkarsh Agrawal, Jimiamma Mafeni Mase, Graziela P Figueredo, Christian Wagner, Mohammad Mesgarpour, and Robert I John. Towards real-time heavy goods vehicle driving behaviour classification in the united kingdom. In 2019 IEEE Intelligent Transportation Systems Conference (ITSC), pages 2330–2336. IEEE, 2019.
- [8] Sara Hooker, Dumitru Erhan, Pieter-Jan Kindermans, and Been Kim. Evaluating feature importance estimates, 2018.
- [9] Olivier Goudet, Diviyani Kalainathan, Philippe Caillou, Isabelle Guyon, David Lopez-Paz, and Michele Sebag, Learning functional causal models with generative neural networks. In *Explainable and interpretable models in computer vision and machine learning*, pages 39–80. Springer, 2018.
- [10] Vaishak Belle. Logic meets probability: Towards explainable ai systems for uncertain worlds. In *International Joint Conference on Artificial Intelligence*, pages 5116–5120, 2017.
- [11] H Chad Lane, Mark G Core, Michael Van Lent, Steve Solomon, and Dave Gomboc. Explainable artificial intelligence for training and tutoring. Technical report, 2005.
- [12] Zachary C Lipton, The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery, *Queue* 16 (3) (2018) 31–57.
- [13] Mark Ryan, In AI We Trust: Ethics, Artificial Intelligence, and Reliability, *Sci. Eng. Ethics* 26 (5) (2020) 2749–2767.
- [14] Andreas Theodorou, Robert H Wortham, Joanna J Bryson, Designing and implementing transparency for real time inspection of autonomous robots, *Connect. Sci.* 29 (3) (2017) 230–241.
- [15] Alexandra Chouldechova, Fair prediction with disparate impact: A study of bias in recidivism prediction instruments, *Big Data* 5 (2) (2017) 153–163.
- [16] Lotfi A Zadeh, Fuzzy sets, *Inf. Control* 8 (3) (1965) 338–353.
- [17] L.-X. Wang, Jerry M Mendel, Generating fuzzy rules by learning from examples, *IEEE Trans. Syst., Man, Cybern.* 22 (6) (1992) 1414–1427.
- [18] Ebrahim H Mamdani, Sedrak Assilian, An experiment in linguistic synthesis with a fuzzy logic controller, *Int. J. Man-Mach. Stud.* 7 (1) (1975) 1–13.
- [19] Leo Breiman, Random forests, *Mach. Learn.* 45 (1) (2001) 5–32.
- [20] Jerome H Friedman, Greedy function approximation: a gradient boosting machine. *Annals of statistics*, 2001, pp. 1189–1232.
- [21] André Altmann, Laura Toloşi, Oliver Sander, Thomas Lengauer, Permutation importance: a corrected feature importance measure, *Bioinformatics* 26 (10) (2010) 1340–1347.
- [22] I. Elizabeth Kumar, Suresh Venkatasubramanian, Carlos Scheidegger, Sorelle Friedler, Problems with shapley-value-based explanations as feature importance measures, in: *International Conference on Machine Learning*, PMLR, 2020, pp. 5491–5500.
- [23] Koen W De Bock, Dirk Van den Poel, Reconciling performance and interpretability in customer churn prediction using ensemble learning based on generalized additive models, *Expert Syst. Appl.* 39 (8) (2012) 6816–6826.
- [24] Bin Xu Zhai, Jianguo Chen, Development of a stacked ensemble model for forecasting and analyzing daily average pm2.5 concentrations in Beijing, China, *Sci. Total Environ.* 635 (2018) 644–658.
- [25] Salome Sanchez, Divish Rengasamy, Christopher J Hyde, Graziela P Figueredo, Benjamin Rothwell, Machine learning to determine the main factors affecting creep rates in laser powder bed fusion, *J. Intell. Manuf.* (2021) 1–21.

[26] Daniel Berrar, Cross-validation, *Encyclopedia of Bioinformatics and Computational Biology* 1 (2019) 542–545.

[27] Jin Gou, Feng Hou, Wenyu Chen, Cheng Wang, Wei Luo, Improving wang-mendel method performance in fuzzy rules generation using the fuzzy c-means clustering algorithm, *Neurocomputing* 151 (2015) 1293–1304.

[28] Ronald R Yager, Dimitar P Filev, Generation of fuzzy rules by mountain clustering, *J. Intell. Fuzzy Syst.* 2 (3) (1994) 209–219.

[29] Myung Won Kim, Joong Geun Lee, and Changwoo Min. Efficient fuzzy rule generation based on fuzzy decision tree for data mining. In FUZZ-IEEE'99. 1999 IEEE International Fuzzy Systems. Conference Proceedings (Cat. No. 99CH36315), volume 3, pages 1223–1228. IEEE, 1999.

[30] Francisco Herrera, Manuel Lozano, Jose Luis Verdegay, Generating fuzzy rules from examples using genetic algorithms, in: *Fuzzy Logic and Soft Computing*, World Scientific, 1995, pp. 11–20.

[31] Diego Alvarez-Estevéz and Vicente Moret-Bonillo. A proposal to solve rule conflicts in the wang-mendel algorithm for fuzzy classification using evidential theory. In 29th Benelux Conference on Artificial Intelligence November 8–9, 2017, Groningen, page 161.

[32] Lotfi A Zadeh, Fuzzy logic and approximate reasoning, *Synthese* 30 (3) (1975) 407–428.

[33] 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, E. Duchesnay, Scikit-learn: Machine learning in Python, *J. Mach. Learn. Res.* 12 (2011) 2825–2830.

[34] Jerome H Friedman, Bogdan E Popescu, et al., Predictive learning via rule ensembles. *Annals of, Appl. Stat.* 2 (3) (2008) 916–954.

[35] David A Winkler and Tu C Le. Performance of deep and shallow neural networks, the universal approximation theorem, activity cliffs, and qsar. *Molecular informatics*, 36(1–2):1600118, 2017.

[36] Oscar Córdón, Francisco Herrera, Pedro Villar, Generating the knowledge base of a fuzzy rule-based system by the genetic learning of the data base, *IEEE Trans. Fuzzy Syst.* 9 (4) (2001) 667–674.



Ben Rothwell is an assistant professor in the department of Mechanical, Materials and Manufacturing Engineering at the University of Nottingham. He received his Ph.D. from the University of Leeds in 2016. From 2016 until 2018, he was a Research Associate with the Gas Turbine and Transmissions Research Centre at the University of Nottingham. Prior to joining the University of Nottingham in 2016, he was an Aerospace instructor at the King Faisal Air Academy in Riyadh, Saudi Arabia. His main research interests include multi physics computational modeling of journal bearings supported by experimental rigs.



Mercedes Torres Torres is Machine Learning Team Leader at B-hive innovations. Her area of expertise is in machine learning and computer vision, with particular emphasis in the development of novel deep learning techniques for small and skewed datasets. Her research portfolio includes multidisciplinary collaborations in areas like Healthcare (QMC), GIS (Dstl), Cybersecurity (Netacea), and Children Protection Services (Nottingham City Council), with funding from EPSRC (Horizon Research Institute, TAS Network, EnLightenUs as Co- I), Innovate UK (KTP, PI), and Huawei (HIRP, PI). She has published in high-impact conferences and journals, such as LREC, AIME, ESWA, FG and IMAVIS.



Divish Rengasamy is a PhD student and a Marie Skłodowska-Curie Early Stage Research Fellow at the Institute of Aerospace Technology in University of Nottingham. He is also currently a machine learning engineer working at a UK-based start-up focusing on LiDAR-based power line inspection. His research focuses on safer and more responsible usage of machine learning in safety-critical applications.



Morgan Alexander is Professor of Biomedical Surfaces in the School of Pharmacy at the University of Nottingham, received his BSc in Materials (1988) and PhD (1992) University of Sheffield. His research into novel biomaterials involves developing materials for application in healthcare, characterising relationships between surfaces and the biological response. Understanding these relationships is critical to the development of next generation biomaterials and it is the theme running through his work across a variety of areas spanning control of bacterial and fungal colonisation, immune cell control for implants, stem engineering cell response with applications in medical devices and cell manufacture.



Jimiama Mafeni Mase is a PhD candidate in Computer Science at the University of Nottingham. His PhD explores the use of expert knowledge in multi-modal decision and information fusion for fair, reliable, interpretable and comprehensive assessments. His expertise includes modelling uncertainty and imprecision in knowledge, information fusion, computer vision, data stream mining, anomaly detection, domain expert engagement and knowledge extraction, development of reliable and fair novel machine learning techniques, and explainable AI. He is currently a member of Computer Vision Lab (CVL) and the Lab for Uncertainty in Data and Decision Making (LUCID) in the University of Nottingham.



David Winkler is Professor of Biochemistry & Chemistry at La Trobe University, Professor of Pharmacy at the University of Nottingham, and Professor of Medicinal Chemistry at Monash University. He applies computational chemistry, AI, and machine learning to the design of drugs, agrochemicals, electro optic materials, nano-materials, and biomaterials. He is ranked 227th of 81,000 medicinal chemists worldwide and has written >250 journal articles and book chapters (4 ISI Highly Cited) and is an inventor on 25 patents.



Aayush Kumar is a consulting manager for Evalueserve, focusing on domains such as information, communication, and technology (ICT). He works with Fortune 500 technology organizations to develop their growth strategies using data science and business analytics. Aayush holds a master's degree in data science from the University of Nottingham. His thesis addressed the development of a Python based Feature Importance Fusion and Interpretation model. Aayush also holds Bachelor of Technology degree in an Electrical & Electronics, from the VIT University in India.



Graziela Figueredo is Assistant Professor in the School of Computer Science and the Research and Innovation Lead at the Digital Research Service at The University of Nottingham. The focus of her research is the development and application of techniques for systems simulation and intelligent data analysis for urban mobility and transportation. Graziela has served as Associate Editor of the IEEE Intelligent Transportation Systems conference and has published over 20 papers in AI for transport research (in journals such as IEEE Transactions on Intelligent Transportation Systems and Accident Analysis and Prevention).