

# The Impact of Feature Reduction Techniques on Defect Prediction Models

Thesis for the degree of Master of Engineering

*Author:*

Masanari Kondo

*Student ID No.:*

17622015

*Chief Supervisor:*

Professor Osamu Mizuno

*Co-Supervisors:*

Professor Teruhisa Hochin

Professor Yu Shibuya

Master's Program of Information Science,  
Graduate School of Science and Technology,  
Kyoto Institute of Technology

February 8, 2019



学位論文内容の要旨（和文）

平成 31 年 2 月 8 日

京都工芸繊維大学大学院  
工芸科学研究科長 殿

工芸科学研究科 情報工学専攻  
平成 29 年入学  
学生番号 17622015  
氏 名 近藤 将成 ㊦

（主任指導教員 水野 修 ㊦）

本学学位規則第 4 条に基づき、下記のとおり学位論文内容の要旨を提出いたします。

1. 論文題目

ソフトウェアメトリクスを用いた不具合予測モデルにおける  
特徴抽出手法の実践的評価

2. 論文内容の要旨（400 字程度）

不具合が含まれたソフトウェアをリリースすると、それに起因した事故を引き起こす可能性がある。そのため、ソフトウェアのリリース前に不具合を全て取り除くことを課題として、様々な手法が研究されている。

この課題を解決する 1 つの手法として、ソースコードを数値化してソースコードの複雑度を計測する基準（メトリクス）を用いて不具合の予測モデルを構築することが行われている。一方で、今日までに多種多様で膨大な量のメトリクスが提案されており、重複するメトリクス（例えば、高い相関がある 2 つのメトリクスのうちの 1 つ）を除去して予測モデルを構築することが必要不可欠となっている。既存研究では、その手段として特徴選択手法が頻繁に用いられ、どの特徴選択手法が良いかを調べる比較評価の研究も行われている。それに対して、我々の研究では、特徴抽出手法を比較評価した。また特徴選択手法とも比較を行い、どちらが優れているのかを調査した。



# **The Impact of Feature Reduction Techniques on Defect Prediction Models**

2019

17622015

*KONDO Masanari*

## **Abstract**

Defect prediction is an important task for preserving software quality. Most prior work on defect prediction uses software features, such as the number of lines of code, to predict whether a file or commit will be defective in the future. So far, many features have been proposed for improving the accuracy of defect prediction. There are several reasons to keep the number of features that are used in a defect prediction model small. For example, using a small number of features avoids the problem of multicollinearity and the so-called ‘curse of dimensionality’. Feature selection and reduction techniques can help to reduce the number of features in a model. Several recent studies have investigated the impact of feature selection techniques on defect prediction. However, there do not exist large-scale studies in which the impact of multiple feature reduction techniques on defect prediction is investigated.

In this paper, we study the impact of eight feature reduction techniques on the performance and the variance in performance of five supervised and five unsupervised defect prediction models. In addition, we compare the impact of the studied feature reduction techniques with the impact of the two best-performing feature selection techniques (according to prior work).



# Index

<b>1. Introduction</b>	<b>1</b>
<b>2. Related work</b>	<b>4</b>
2.1 Defect prediction . . . . .	4
2.2 Feature selection and reduction . . . . .	5
<b>3. Methodology</b>	<b>9</b>
3.1 Studied datasets . . . . .	9
3.2 Studied defect prediction models . . . . .	9
3.3 Studied feature reduction techniques . . . . .	15
3.3.1 Traditional feature reduction techniques . . . . .	15
3.3.2 Neural network-based feature reduction techniques . . . . .	18
3.4 Studied feature selection techniques . . . . .	18
3.5 Area under the receiver operating characteristic curve (AUC) . . . . .	19
3.6 Preprocessing . . . . .	19
3.6.1 Preprocessing for traditional feature reduction techniques . . . . .	20
3.6.2 Preprocessing for neural network-based feature reduction techniques . . . . .	20
3.7 Out-of-sample bootstrap sampling . . . . .	20
<b>4. Experimental setup</b>	<b>22</b>
<b>5. Results</b>	<b>24</b>
5.1 <b>RQ1: What is the impact of feature reduction techniques on the performance of defect prediction models?</b> . . . . .	24
5.2 <b>RQ2: What is the impact of feature reduction techniques on the variance of the performance across defect prediction models?</b> . . . . .	30
5.3 <b>RQ3: How do feature selection techniques compare to feature reduction techniques when applied to defect prediction?</b> . . . . .	37
5.3.1 Why do feature reduction techniques work well in the AEEEM dataset?	42

5.3.2	Comparing feature selection and reduction techniques along the dimensions of understandability and execution time . . . . .	45
<b>6.</b>	<b>Discussion: Which features are generated by the feature reduction techniques?</b>	<b>46</b>
<b>7.</b>	<b>Threats to validity</b>	<b>52</b>
7.1	External validity . . . . .	52
7.2	Internal validity . . . . .	52
<b>8.</b>	<b>Conclusion</b>	<b>53</b>
	<b>Acknowledgment</b>	<b>55</b>
	<b>References</b>	<b>56</b>



# 1. Introduction

Software developers have limited time to test their software. Hence, developers need to be selective when deciding where to focus their testing effort. Defect prediction models help developers focus their limited testing effort on components that are the most likely to be defective. Because detecting defects at an early stage can considerably reduce the development cost [1], defect prediction models have received widespread attention in software engineering research.

Many software features (e.g., software complexity features) can be used in defect prediction models [2–6]. However, it is important to carefully select the set of features that is used in such models, as using a set of features that is too large does not automatically result in better defect prediction. For example, prior studies showed that reducing the number of features avoids the problem of multicollinearity [7] and the curse of dimensionality [8]. Hence, many of the existing defect prediction models used feature selection or reduction techniques [6, 9–28].

Feature selection techniques reduce the number of features in a model by selecting the most important ones, while feature reduction techniques reduce the number of features by creating new, combined features from the original features. Recent studies [9, 18] investigated the impact of feature *selection* techniques on the performance of defect prediction models. However, this paper describes the first large-scale study on multiple feature *reduction* techniques and their impact on a large number of prediction models.

In this paper, we compared the impact of the original features, features that are generated using *traditional feature reduction techniques* (i.e., PCA [6], FastMap [29], feature agglomeration [30], random projections [31], TCA [32] and TCA+ [28]), and features that are generated using *neural network-based feature reduction techniques* (i.e., restricted Boltzmann machine (RBM) [33] and autoencoder (AE) [34]) on defect prediction models. In addition, we compared the impact of features that are generated using feature reduction techniques with features that are selected using the best-performing feature selection techniques in prior work (correlation and consistency-based feature selection) [9, 18]. We compared the features along two dimensions: their performance (area under the receiver operating characteristic curve (AUC)), and their performance variance (interquartile range (IQR)). The receiver operating characteristic curve is created by plotting the true positive rate against the false positive rate, hence, the

AUC represents the balance between the true positive rate and the false positive rate. The IQR represents the variance of a data distribution.

We conducted experiments on three publicly available datasets that contain software defects (the PROMISE [14], cleaned NASA [35], and AEEEM [6] datasets). Our ultimate goal is to identify which feature reduction techniques yield new, powerful features that preserve the predictive power of the original features, and improve the prediction performance compared to feature selection techniques. We studied the impact of feature reduction techniques on five supervised learning and five unsupervised learning models for defect prediction in our experiments. In particular, we focus on the following research questions:

**RQ1: What is the impact of feature reduction techniques on the performance of defect prediction models?**

*Motivation:* Reducing the number of features in a model can address the multicollinearity problem [7] and the curse of dimensionality [8]. In this RQ, we studied how feature reduction techniques impact the performance of supervised and unsupervised defect prediction models.

*Results:* Feature agglomeration and TCA can reduce the number of features, while preserving an AUC that is as good as that of the original features for supervised models. In addition, the AUC of unsupervised defect prediction models is significantly better when preprocessing the features with neural network-based feature reduction techniques than other feature reduction techniques.

**RQ2: What is the impact of feature reduction techniques on the variance of the performance across defect prediction models?**

*Motivation:* The AUC for a dataset can vary across defect prediction models. Hence, it can be challenging for practitioners to choose the defect prediction model that performs best on their data. If all defect prediction models have a small performance variance for a dataset, practitioners can avoid having to make this challenging choice.

*Results:* Neural network-based feature reduction techniques (RBM and AE) generate features that improve the variance of the performance across different defect prediction models in many cases when used in a supervised or unsupervised defect prediction

model. In addition, almost all feature reduction techniques (except PCA) generate features that improve the variance of the performance across different defect prediction models in many cases when used in an unsupervised defect prediction model.

**RQ3: How do feature selection techniques compare to feature reduction techniques when applied to defect prediction?**

*Motivation:* Prior work [9, 18] showed that several feature selection techniques outperform the original models. In this RQ, we studied how feature selection techniques compare to feature reduction techniques.

*Results:* For the supervised defect prediction models, the studied feature selection techniques (correlation and consistency-based feature selection) outperform all the studied feature reduction techniques. However, for the unsupervised defect prediction models, the neural network-based feature reduction techniques (RBM and AE) outperform the other studied feature selection/reduction techniques.

Our results provide practitioners with advice on which feature selection/reduction technique to use in combination with a defect prediction model. In particular, we recommend to use a feature selection technique when using a supervised defect prediction model, and a neural network-based feature reduction technique when using an unsupervised defect prediction model, as these feature selection/reduction techniques improve the variance across defect prediction models, while outperforming the other feature reduction techniques.

The organization of our paper is as follows. Section 2 introduces related work. Section 3 presents our methodology. Section 4 presents experimental setup. Section 5 presents the results of our experiment. Section 6 discusses these results. Section 7 describes the threats to the validity of our findings. Section 8 presents the conclusion.

This work was accepted on Empirical Software Engineering [36].

## 2. Related work

In this section, we discuss related work on defect prediction, and feature selection and reduction.

### 2.1 Defect prediction

Defect prediction approaches can be divided in two categories: approaches that use supervised learning and unsupervised learning. Supervised defect prediction models need training data and test data. Usually, the training data and test data are collected from the same, or very similar projects from the same organization (within-project defect prediction). However, sometimes it is difficult to collect sufficient training data from the same project or organization (e.g., when the project is new).

In such cases, cross-project defect prediction can be a viable alternative solution. Cross-project defect prediction uses training datasets that contain data that comes from multiple projects, or even from multiple organizations. Cross-project defect prediction offers a solution for the problem of within-project defect prediction since small datasets can be extended with data from other projects. Hence, cross-project defect prediction is one of the important research themes in defect prediction research. However, cross-project defect prediction has several challenges [13, 17, 37–39]. For instance, cross-project defect prediction has the problem of heterogeneous original features across datasets [17]. There exists one solution for this challenge [40], however; converging datasets collected from multiple companies to one dataset when they contain different features still remains an open challenge and further investigation is needed.

One way to overcome the challenge of heterogeneity is to use unsupervised defect prediction models [16, 41–45]. Unsupervised models have the advantage that they do not need a training dataset, and therefore, are not affected by the problem of heterogeneous features [16, 41–45]. Recently, Zhang et al. [41] summarized the accuracy of several supervised and unsupervised models for defect prediction. They concluded that connectivity-based unsupervised models have an accuracy that is as good as that of supervised models. Therefore, unsupervised models

are a viable alternative to supervised models for defect prediction. Note that unsupervised models do not need training data, and are therefore always within-project defect prediction models.

## 2.2 Feature selection and reduction

Using feature selection or reduction technique has the advantage of addressing the *curse of dimensionality* [8]. This problem is originally considered a dynamic optimization problem. However, machine learning models also need to consider this problem. The problem occurs when having a large number of features yet a small sample size in machine learning models. In this case, the sample size is not large enough to search the representations of all the combinations of features, and to generalize their parameters (resulting in overfitting) [49]. Hence, the prediction performance of these models for unseen data would be worse, and could lead to a classification error. Feature selection or reduction techniques can address this problem by removing or combining redundant and irrelevant features. In addition, prior work [46, 50] reported that the prediction performance of a model is often determined by only a few features. Therefore, feature selection or reduction techniques can often be applied without negatively affecting the prediction performance.

In this paper, we define feature selection and reduction as follows:

- Feature selection techniques reduce the number of features by selecting a subset of the original features.
- Feature reduction techniques reduce the number of features by combining original features into new features.

Several researchers studied the impact of feature selection techniques on defect prediction models [9–22, 24–26, 46]. Table 2.1 gives an overview of prior work that uses feature selection or reduction techniques. For instance, Ghotra et al. [9] summarized the impact of feature selection techniques for defect prediction. They reported that correlation-based feature selection outperforms the other feature selection techniques. In addition, they showed that the impact of feature selection techniques varies across the studied datasets. Xu et al. [18] also summarized the impact of feature selection techniques on defect prediction models. They also reported that

**Table 2.1 Overview of prior work that uses feature selection or reduction techniques in combination with defect prediction.**

<b>Technique</b>	<b>Family</b>	<b>Example of Techniques</b>	<b>References</b>	
Feature Selection	Filter-based Feature Ranking	Chi-Square Correlation	[9, 11, 12, 14, 18, 20, 21, 24, 26, 46]	
	Filter-based Subset Selection	Correlation-based Feature Subset Selection	[9, 10, 18, 20, 21, 24]	
	Wrapper-based Subset Selection	Logistic Regression	[9, 18, 21, 22, 25, 26]	
	Greedy-based Feature Selection	Greedy Forward Selection Algorithm	[24]	
	Distance-based Feature Selection	EM Algorithm	[13]	
	State-of-the-Art	MVS	[16]	
	Others	Significance	[15, 17, 19, 26]	
	Feature Reduction	PCA	PCA	[6, 9, 10, 18, 21, 27, 38, 47, 48]
		State-of-the-Art	TCA	[27, 28]

the effectiveness of feature selection techniques exhibits significant differences across studied datasets. Menzies et al. [46] used the information gain measure to rank features based on their importance in a defect prediction model. They found that defect prediction can often be done using only a small set of features. For example, they reduced the number of features in the CM1 project from 38 to 3 without affecting the prediction performance.

The impact of feature reduction techniques on defect prediction models has not been studied as extensively. Most researchers use principal component analysis (PCA) [6,9,10,18,21,27,38,47,48], and only a few researchers use other feature reduction techniques [27,28]. For instance, D'Ambros et al. [6] compared class-level defect prediction models to present a benchmark for defect prediction. They used PCA to avoid the problem of multicollinearity [7] among the independent variables. Nagappan et al. [47] predicted post-release defects using complexity features. They built prediction models using PCA to avoid the problem of multicollinearity. Neumann [48] proposed PCA-ANN which is a combination of PCA and artificial neural networks. Neumann also used PCA to avoid the problem of multicollinearity. Challagulla et al. [10] compared several prediction models for identifying defects. In addition, they compared PCA with feature selection techniques such as feature subset selection. They concluded that feature selection techniques are better than PCA, and that PCA did not add any advantages for defect prediction. Rathore et al. [21] compared the performance of feature selection techniques. In this comparison, they also used PCA. They found that PCA is one of the best-performing techniques in this comparison. Ren et al. [27] extended PCA to address class imbalance problem for defect prediction. Nam et al. [28] applied transfer component analyses (TCA and TCA+) to training and test data to convert the data to be closer than the original data. This process addressed the challenge of heterogeneity of training data and test data in cross-project defect prediction. Nam's approach significantly improved cross-project defect prediction performance.

Despite the amount of prior work on feature reduction techniques (mostly on PCA), no prior work has conducted a systematic study of the impact of feature reduction techniques on defect prediction models. In this paper, we provide such a systematic study of the impact of eight feature reduction and two best-performing feature selection techniques on five supervised and five unsupervised defect prediction models.

Finally, Peters et al. [51] showed that often not only the features can be reduced without

negatively affecting the performance, but also the amount of the rows in the data.



## 3. Methodology

In this section, we describe our methodology. In particular, we discuss our studied datasets, defect prediction models, feature selection techniques, feature reduction techniques, evaluation measure, our preprocessing steps, and our validation scheme.

### 3.1 Studied datasets

In our work, we used three publicly available datasets (the PROMISE [14], cleaned NASA [35] and AEEEM [6] datasets) that were used by prior work [41] on supervised and unsupervised defect prediction models. Table 3.1 describes the studied datasets. All datasets contain popular software features for measuring source code complexity (see Table 3.2 for a summary of the used features). Each entity in a dataset is labelled as defective or clean.

The PROMISE dataset contains several types of projects. We chose the 10 projects that were used by prior work [41], to ease the comparison of our results with prior work. All studied PROMISE projects have the same number of features. The PROMISE dataset contains the Chidamber and Kemerer (CK) features [52] and an additional set of object-oriented (OO) features [2].

The NASA dataset [35] contains 11 projects. Each project in the NASA dataset has a different number of features. The NASA dataset contains McCabe features [53] and Halstead features [54]. We used the cleaned version [35] of the NASA dataset, because prior studies [35,55] showed that the original version of the NASA dataset contains inconsistent and mislabeled data.

The AEEEM dataset [6] contains five projects. All projects have the same number of features. Like the PROMISE dataset, the AEEEM dataset contains the CK and OO features. However, the AEEEM data also contains the number of previous defects [3], change features [4], complexity code change features [5], and the churn and entropy of the CK and OO [6] features.

### 3.2 Studied defect prediction models

We focused on defect prediction models that were used by prior work [41], to make our results easier to compare. We studied five supervised models and five unsupervised models.

**Table 3.1 Description of studied projects**

Studied Dataset	Project	# of	# of	% Defective	# of	# of
		Entities	Defective		Features*	Studied Features*
PROMISE	Ant v1.7	745	166	22.3	24	20
	Camel v1.6	965	188	19.5	24	20
	Ivy v1.4	241	16	6.6	24	20
	Jedit v4.0	306	75	24.5	24	20
	Log4j v1.0	135	34	25.2	24	20
	Lucene v2.4	340	203	59.7	24	20
	POI v3.0	442	281	63.6	24	20
	Tomcat v6.0	858	77	9.0	24	20
	Xalan v2.6	885	411	46.4	24	20
	Xerces v1.3	453	69	15.2	24	20
NASA	CM1	327	42	12.8	38	37
	JM1	7,782	1,672	21.5	22	21
	KC3	194	36	18.6	40	39
	MC1	1,988	46	2.3	39	38
	MC2	125	44	35.2	40	39
	MW1	253	27	10.7	38	37
	PC1	705	61	8.7	38	37
	PC2	745	16	2.1	37	36
	PC3	1,077	134	12.4	38	37
	PC4	1,287	177	13.8	38	37
	PC5	1,711	471	27.5	39	38
AEEEM	Eclipse JDT Core	997	206	20.7	291	212
	Equinox	324	129	39.8	291	212
	Apache Lucene	691	64	9.3	291	212
	Mylyn	1,862	245	13.2	291	212
	Eclipse PDE UI	1,497	209	14.0	291	212

\* We removed features that are not related to source code. For instance, the name of the file, name of the class and the version. Hence, the number of studied features are different from the total number of features.

**Table 3.2 Studied features**

<b>Studied Dataset</b>	<b>Features</b>	<b>Reference</b>
PROMISE	CK	Chidamber et al. [52]
	OO	Basili et al. [2]
NASA	McCabe	McCabe [53]
	Halstead	Halstead [54]
AEEEM	CK	Chidamber et al. [52]
	OO	Basili et al. [2]
	number of previous defects	Kim et al. [3]
	change features	Moser et al. [4]
	complexity code change features	Hassan [5]
	churn of CK and OO	D'Ambros et al. [6]
entropy of CK and OO	D'Ambros et al. [6]	

Below we give a brief overview of the ideas behinds these models. For a detailed overview, we refer the reader to the original papers in which these models were introduced. We studied the following supervised defect prediction models:

- *Logistic Regression (LR)* [56]: LR is one of the most commonly used models for defect prediction. LR expresses the relationship between one or more independent variables (i.e., the original features) and one dependent variable (i.e., defective or clean) using a polynomial expression and a *sigmoid function* [57].
- *Decision Tree (J48)* [58]: J48 is a decision tree implementation in WEKA [59]. The decision tree uses a tree structure to decide on the dependent variable. In this tree, each node corresponds to one of the independent variables with a condition. J48 traverses the tree from the root to a leaf, while taking into account the input entity and the conditions in the tree. Each leaf corresponds to a value of the dependent variable.
- *Random Forest (RF)* [60]: RF is a popular *ensemble learning* model. RF builds multiple decision trees based on subsets of training data that are randomly selected. RF decides on a value of the dependent variable by taking the result of a majority of the decision trees.
- *Naive Bayes (NB)* [61]: NB is a probability-based classifier that follows Bayes' theorem. Bayes' theorem describes the probability of an event, given knowledge of conditions that could be related to the event.
- *Logistic Model Tree (LMT)* [62]: LMT is a classifier which combines a decision tree and a logistic regression model. Like the decision tree, LMT follows a tree structure. However, LMT uses logistic regressions instead of values in the leaves.

We used the caret library in R [63] to optimize the parameters of the supervised models as suggested by Tantithamthavorn et al. [64].

We studied the following unsupervised defect prediction models:

- *Spectral Clustering (SC)* [65]: SC labels entities using a graph that is based on similarities across entities. In this graph, each node is an entity and each edge represents the similarity of the entities it connects. SC cuts sparse edges in this graph by classifying eigenvectors of the *Laplacian matrix* [65] of the graph. Following this process, SC divides the entities into two groups.

- *k-means (KM) [66]*: KM is a popular clustering approach. KM classifies entities based on the distances between entities and the center of a class (i.e., the mean of all entities in that class). In this paper, we used the Euclidean distance as the distance metric.
- *Partition Around Medoids (PAM) [67]*: PAM is an approach that is similar to KM. While KM uses the center of a class, PAM uses *medoids*. A medoid is an entity of which the sum of all the distances to the other entities in a class is at its minimum. Because PAM uses the medoid instead of the center, PAM is less likely to be affected by outliers than KM.
- *Fuzzy C-Means (FCM) [68]*: FCM is also an approach that is similar to KM. While KM classifies each entity to only one class in its process, FCM allows entities to be a member of more than one class. The membership is expressed as a probability.
- *Neural Gas (NG) [69]*: NG is an approach that is similar to a *self-organizing map* [70]. NG generates *weighted points* which have random features. Hence, the weighted points are distributed across the feature space. For each learning iteration, the features of the weighted points are updated by distances to closer entities. Finally, the weighted points become the class centers.

We used the default parameters of the implementations for the unsupervised models. We set the number of clusters to two, as defect prediction is a binary problem. Table 3.3 shows the libraries that we used for the implementation of the defect prediction models.

**Labeling technique in the unsupervised models:** The unsupervised models classify the data in two unlabeled clusters. We adopted the following heuristic to identify the defective cluster: “*For most features, software entities containing defects generally have larger values than software entities without defects*” [41]. In particular, we used the sum of row average of the normalized features in each cluster, to decide which cluster contains the defects [41]. To calculate the sum of row average, we first summed the entity values in each cluster, respectively. Then, we calculated the average values for each cluster. The cluster with the larger average value was identified as the cluster with the defective entities.

**Table 3.3 Libraries used for experiments**

<b>Defect Prediction Models</b>	<b>Libraries</b>
LR	The caret library in R
RF	The caret library in R
NB	The caret library in R
J48	The caret library in R
LMT	The caret library in R
SC	Zhang et al.'s implementation [41] (see app. A)
KM	The cclust library in R
PAM	The cluster library in R
FCM	The e1071 library in R
NG	The cclust library in R

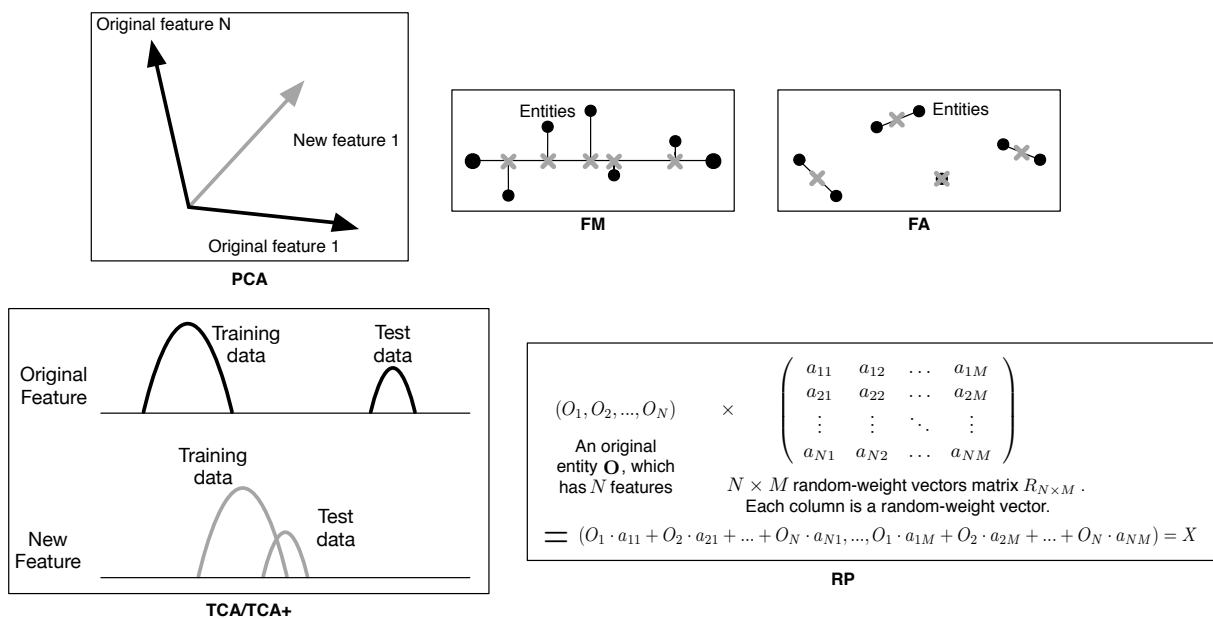
### 3.3 Studied feature reduction techniques

In this subsection, we discuss the studied feature reduction techniques. We studied two types of feature reduction techniques: *traditional* and *neural network-based* feature reduction techniques. We give a brief overview of the core concepts of each feature reduction technique. For more precise details, we refer to the references that are mentioned for each technique. Figure 3.1 shows a visualization of the traditional feature reduction techniques (PCA, FM, FA, TCA/TCA+, RP). Figure 3.2 gives an overview of neural network-based feature reduction techniques (RBM and AE).

#### 3.3.1 Traditional feature reduction techniques

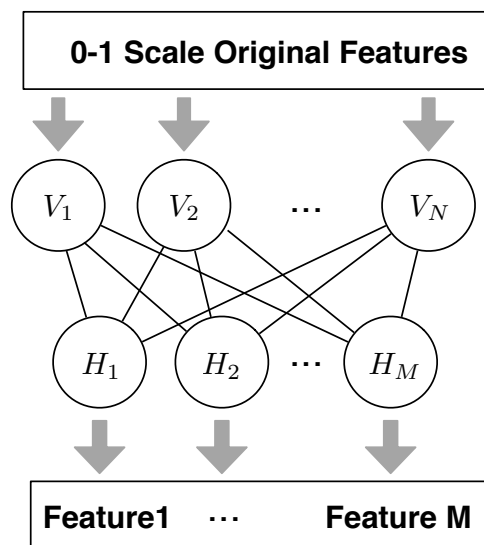
We studied the following traditional feature reduction techniques.

- *Principal Component Analysis (PCA)*: PCA is one of the most commonly used feature reduction techniques in defect prediction [6, 9, 10, 18, 21, 23, 27, 38]. PCA reduces the number of features by projecting the original set of features on a smaller number of principal components.
- *FastMap (FM)*: For  $N$  original features, FM [29] first generates a  $(N-1)$ -dimensional orthogonal hyper-plane of the line between two entities that are far from each other. Second, FM projects the other entities on this hyper-plane. Because FM projects the entities on the  $N - 1$  orthogonal hyper-plane, we can reduce one feature from the original features. FM repeats this procedure until we get the required number of new features. For instance, if we want three features to visualize our data from the  $N$  original features, we repeat the procedure  $N-3$  times.
- *Feature Agglomeration (FA)*: FA is a simple hierarchical clustering algorithm [30]. FA starts by creating a new feature from each original feature. Then, FA merges the two nearest (based on their Euclidean distance) features into one feature, and repeats this process until the desired number of features is reached.
- *Transfer Component Analysis (TCA and TCA+)*: TCA [32] creates new features from the original features by projecting them on so-called transfer components (similar to PCA). However, the goal of TCA is not to reduce the number of features, but to reduce the gap



**Figure 3.1** A visual overview of the core concepts of the traditional feature reduction techniques. The black symbols represent the original features (or entities in FM and FA) and the gray symbols represent the newly-generated features (or entities in FM and FA). RP (Random projection) transforms an original entity to a new entity using  $M$  random-weight vectors.





**Figure 3.2** An overview of neural network-based feature reduction techniques (RBM and AE). RBM and AE convert the original features ( $V_i$ ), which values must range between 0 and 1, into  $M$  new features ( $H_i$ ). Note that the original input data may need to be preprocessed to satisfy the 0-1 range requirement.

between the distribution of the training and testing data. During this process, the number of features is often reduced. Hence, TCA can be used as a feature reduction technique. TCA+ [28] is an extension of TCA, which optimizes the data using a preprocessing step according to the gap between the distribution of the training and testing data, such as scaling the original features between 0 and 1 instead of using the  $z$ -score.

- *Random Projection (RP)*: RP projects the original  $N$ -dimensional features onto  $M$  generated features ( $M \ll N$ ) using a  $N \times M$  random-weight vectors matrix [31]. The equation of RP is as follows:

$$\mathbf{X} = \mathbf{O} \times R_{N \times M}$$

where  $\mathbf{X}$  is a generated  $M$ -dimensional vector entity,  $\mathbf{O} = (O_1, O_2, \dots, O_N)$  is an original entity, and  $R_{N \times M}$  is a random-weight vectors matrix. For example, if we want three features from  $N$  original features, we prepare three random-weight vectors with  $N$  random values in each of them. The random values are selected such that  $\mathbf{X}$  represents the original features.

### 3.3.2 Neural network-based feature reduction techniques

We studied the following neural network-based feature reduction techniques.

- *Restricted Boltzmann Machine (RBM)*: An RBM [33] automatically extracts important information from the original features as weights and biases on a two-layered neural network. Each node in the first-layer corresponds to an original feature, and each node in the second-layer corresponds to a new feature. We use the output of the second-layer as the new features.
- *Autoencoder (AE)*: AE [34] and RBM are similar, but trained differently. In RBM, the network is trained based on a probability distribution. In AE, the network is trained using the difference between the original and the generated features.

## 3.4 Studied feature selection techniques

We studied the correlation-based (CFS) and consistency-based feature selection techniques (ConFS). These techniques were reported as the best-performing feature selection techniques

in prior studies [9, 18]. Below, we give a brief overview of these techniques.

- Correlation-based feature selection (CFS) [71]: CFS selects a subset of features based on their *correlation*. The selected features have strong correlations with the class label (clean or defective), while having a low correlation with each other.
- Consistency-based feature selection (ConFS) [72]: ConFS uses the *consistency* of the class label across the entities instead of the correlation. For example, if file A has a feature set (10, 20, 40, defective) and file B has a feature set (10, 20, 30, clean), we can identify the defective and clean entities using the third feature. However, if a feature reduction technique removes the third feature, file A and file B have the same feature set except for the class label. In that case, these entities are inconsistent. Using this information, ConFS selects the best feature subset from the original features.

### 3.5 Area under the receiver operating characteristic curve (AUC)

We used the Area Under the receiver operating characteristic Curve (AUC) as the performance measure since AUC is not affected by the skewness of defect data [41, 73]. The receiver operating characteristic (ROC) curve is created by plotting the false positive rate (on the x-axis) and the true positive rate (on the y-axis) at various thresholds. In our experiment, the false positive rate is defined as the portion of clean entities that are identified as defective; the true positive is defined as the portion of defective entities that are identified as defective. The threshold is used to label an entity as clean or defective by checking whether its predicted probability is over the threshold. The AUC is the area under the ROC curve. The values of the AUC range between 0 and 1; a perfect classifier has an AUC of 1, while a random classifier has an AUC of 0.5.

### 3.6 Preprocessing

Most feature reduction techniques require the data to be preprocessed. We detail the preprocessing step below.

### 3.6.1 Preprocessing for traditional feature reduction techniques

The traditional feature reduction techniques require features that are normalized to a mean of 0 and a variance of 1 using the  $z$ -score [41]. The  $z$ -score is calculated as follows:

$$\mathbf{X}_z = \frac{\mathbf{X}_{org} - \mu}{\sigma} \quad (3.1)$$

where  $\mu$  is a mean of the value of the feature for all entities and  $\sigma$  is the standard deviation of the value of the feature for all entities.

### 3.6.2 Preprocessing for neural network-based feature reduction techniques

The neural network-based feature reduction techniques require either binary features or features that are between 0 and 1. Hence, we scale the original features as follows:

$$\mathbf{X}_{scaled} = \frac{\mathbf{X}_{org} - X_{min}}{X_{max} - X_{min}} \quad (3.2)$$

where  $\mathbf{X}_{org}$  is a vector of the value of a particular feature for all entities.  $X_{min}$  is the smallest value of the feature and  $X_{max}$  is the largest value of the feature for all entities [74].

## 3.7 Out-of-sample bootstrap sampling

Bootstrap sampling is a validation technique that is used to estimate the performance of a model for unseen data. The technique is based on random sampling with replacement. *Out-of-sample bootstrap sampling* is a bootstrap sampling technique that estimates the future performance of a defect prediction model more accurately than a cross-validation scheme [73,75]. Hence, we used the out-of-sample bootstrap sampling technique instead of a conventional validation technique such as 10-fold cross-validation. The process of the out-of-sample bootstrap sampling is as follows:

1. Sample  $N$  data points following the distribution of the original dataset, with  $N$  data points, while allowing for replacement.
2. Train a model using the sampled  $N$  data points, and test it using the data points that were not sampled.
3. Repeat steps 1 and 2  $M$  times.

4. Report the average/median performance as the performance estimate.

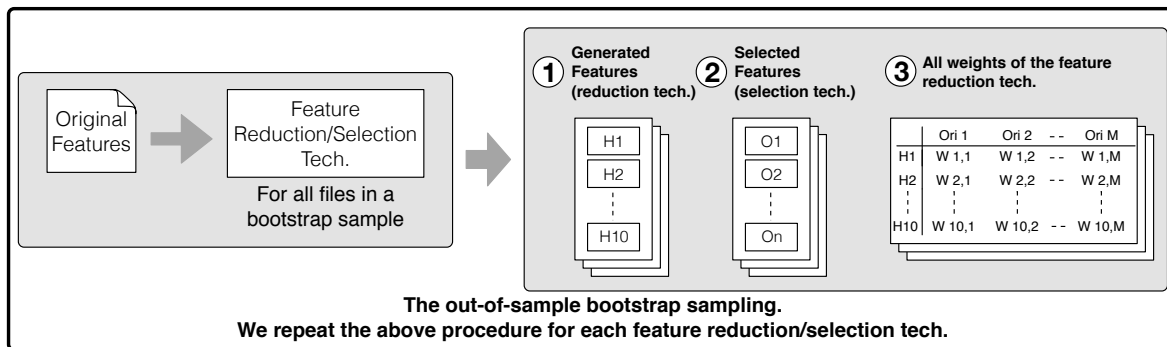
We used the out-of-sample bootstrap sampling under the condition where  $M = 100$  and we report the median performance.

## 4. Experimental setup

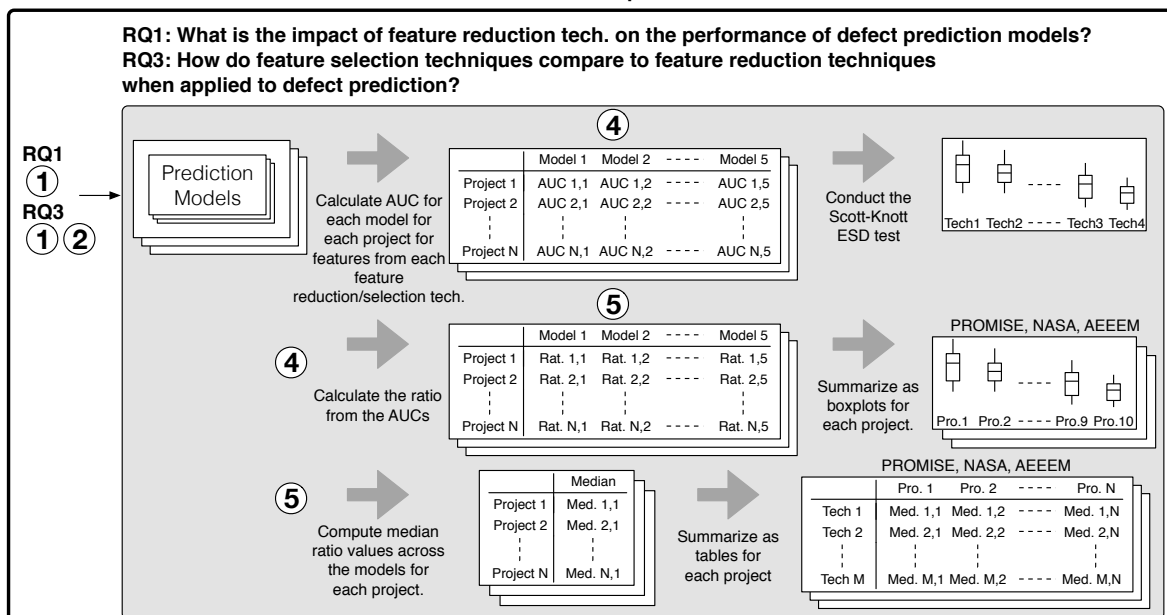
In this section, we give an overview of the setup of our experiments. The results are presented in Section 5. Figure 4.1 shows the steps of our experiments. We first conducted the out-of-sample bootstrap sampling on our studied datasets to generate and select features using each of the studied feature reduction and selection techniques. We then preprocessed the original features of each *bootstrap sample* as described in Section 3.6. We generated eight new feature sets (one for each feature reduction technique) for each bootstrap sample. Hence, we generated 800 new feature sets using feature reduction in total. Furthermore, the two studied feature selection techniques selected two feature subsets (one for each feature selection technique) for each bootstrap sample. Hence, we selected 200 feature subsets using feature selection in total.

The smallest number of features in the studied datasets is 20 (i.e., in the PROMISE dataset). Hence, to be able to observe the impact of a feature reduction technique, we configured each feature reduction technique to generate 10 features (H1–H10). However, PCA uses variance to decide on the number of generated features [6, 9]. Therefore, each bootstrap sample results in a different number of generated features using PCA. We configured PCA to retain 95% of the variance in the data [6, 9]. The median number of generated features by PCA in our experiments was 12 in the PROMISE dataset, 10 in the NASA dataset and 34 in the AEEEM dataset.

The experimental setup for each RQ is discussed in the next section.

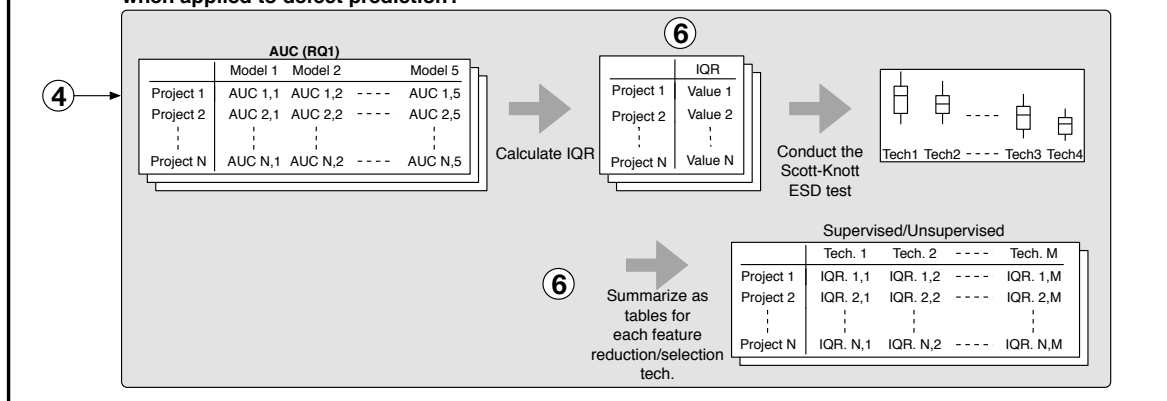


**First Step**




**RQ2: What is the impact of feature reduction tech. on the variance of the performance across defect prediction models?**

**RQ3: How do feature selection techniques compare to feature reduction techniques when applied to defect prediction?**



**Second Step**



**Third Step**

**Figure 4.1 Overview of our experimental design.**

## 5. Results

In this section, we present the results of our experiments. For each RQ, we discuss the motivation, approach and results.

### 5.1 RQ1: What is the impact of feature reduction techniques on the performance of defect prediction models?

*Motivation:* Reducing the number of features that are used in a defect prediction model can be beneficial for addressing the curse of dimensionality and multicollinearity of the model. There exist two ways to reduce the number of features in a model: (1) by *selecting* the most important features, and (2) by *reducing* the number of features by creating new, combined features from the original features. Prior work has systematically studied the impact of feature selection techniques on defect prediction [9, 18], but no work has conducted a large-scale study of the impact of feature reduction techniques on defect prediction. Hence, in this RQ, we studied the impact of feature reduction techniques on the performance (AUC) of defect prediction models.

*Approach:* We used each feature set that was generated by a feature reduction technique as input to the studied 5 supervised and 5 unsupervised defect prediction models. We used the AUC as the performance measure. Because we calculated the AUC of a defect prediction model using the out-of-sample bootstrap sampling 100 times for each feature reduction technique, each model has 100 AUC values. Hence, we used the median value to represent the median performance of a defect prediction model using a certain feature reduction technique. Because we studied 26 projects, our experiments yielded 260 median AUC values for each feature reduction technique (5 supervised models\*26 projects+5 unsupervised models\*26 projects). For comparison, we also calculated the performance of the studied defect prediction models without applying a feature reduction technique (indicated as *ORG*). Note that we did within-project defect prediction in our experiments.

We used the Scott-Knott ESD test [75] (using a 95% significance level) to compare the median AUC values across feature reduction techniques. The Scott-Knott test is a hierarchical clustering algorithm that ranks the distributions of values. In particular, distributions that are



not statistically significantly different are placed in the same rank. The Scott-Knott ESD test is an extension of the Scott-Knott test, which not only ranks based on significance, but also on Cohen's  $d$  effect size [76]. The Scott-Knott ESD test places distributions which are not significantly different, or have a negligible effect size, in the same rank. We used the `ScottKnottESD` R package<sup>(\*)</sup> that was provided by Tantithamthavorn [64].

**Project-level analysis:** the aforementioned procedure combines the results of all projects. However, this procedure prevents us from understanding differences for each project. Hence, we also studied the performance at the project-level.

We compared the ratios of the AUCs (the median AUCs across all bootstrap samples) of each feature reduction technique to the original models. We calculated this ratio as follows:

$$\text{ratio} = \frac{AUC_{FR}}{AUC_{ORG}}$$


Where  $AUC_{ORG}$  is the AUC of a prediction model using the original features, and  $AUC_{FR}$  is the AUC of a prediction model using the features that were generated by a particular feature reduction technique. Hence, a ratio larger than 1 indicates that the feature reduction technique improved the AUC compared to the original models, while a ratio smaller than 1 indicates that the feature reduction technique reduced the AUC. We computed the median ratio across the five studied supervised and unsupervised prediction models.

We used the aforementioned ratio to analyze performance at the project-level. The project-level analysis shows the impact of the different feature reduction techniques in every single project and dataset. Figure 5.3 shows the distributions of the ratios for each studied project for the supervised and unsupervised prediction models, respectively. Each boxplot contains 40 ratio values (5 prediction models \* 8 feature reduction techniques). In addition, we show the median ratios for the best-performing feature reduction techniques as tables for deeper analysis (Table 5.1 and Table 5.2). These median ratios were computed from five AUC values (one for each studied prediction model).

**Results: FA and TCA can preserve the performance of the original defect prediction models, while at the same time reducing the number of features.** Figure 5.1(a) and Figure 5.1(b) show that the performance of the supervised and unsupervised defect prediction

---

(\*)1): <https://github.com/klainfo/ScottKnottESD>



(a) The supervised models (b) The unsupervised models

**Figure 5.1 The Scott-Knott ESD test results for the supervised (logistic regression, random forest, naive Bayes, decision tree, and logistic model tree) and the unsupervised (spectral clustering,  $k$ -means, partition around medoids, fuzzy C-means, neural-gas) models. We use two colors (dark gray and light gray) and two lines (solid and dashed lines) to represent the ranks according to the Scott-Knott ESD test. The adjacent boxplots with the same color and line indicate the same rank. Otherwise, the rank is changed at that point. The different rank indicates a statistical significant difference with small effect size or over. Each boxplot has 130 median AUC values (5 defect prediction models times 26 projects). The x-axis refers to the feature reduction techniques; the y-axis refers to the AUC values.**

models does not decrease when applying FA or TCA. Hence, these feature reduction techniques can safely be applied to reap the benefits of a reduced number of features. In particular, FA and TCA work well for supervised models. Interestingly, the performance of the supervised and unsupervised defect prediction models is significantly lower when using TCA+ (which is an extension of TCA), compared to the original TCA.


**The neural network-based feature reduction techniques (RBM and AE), significantly outperform traditional feature reduction techniques for the unsupervised defect prediction models.** Figure 5.1(b) shows the AUC values and the results of the Scott-Knott ESD test for the unsupervised models after applying the studied feature reduction techniques.

The highest rank contains only the two studied neural network-based techniques: RBM and AE. Hence, the neural network-based feature reduction techniques can significantly improve the AUC compared to the original models and other feature reduction techniques. However, these neural network-based feature reduction techniques do not outperform ORG for the supervised models. In Section 6 we further investigate why neural network-based feature reduction techniques work well for the unsupervised, but not for the supervised defect prediction models.

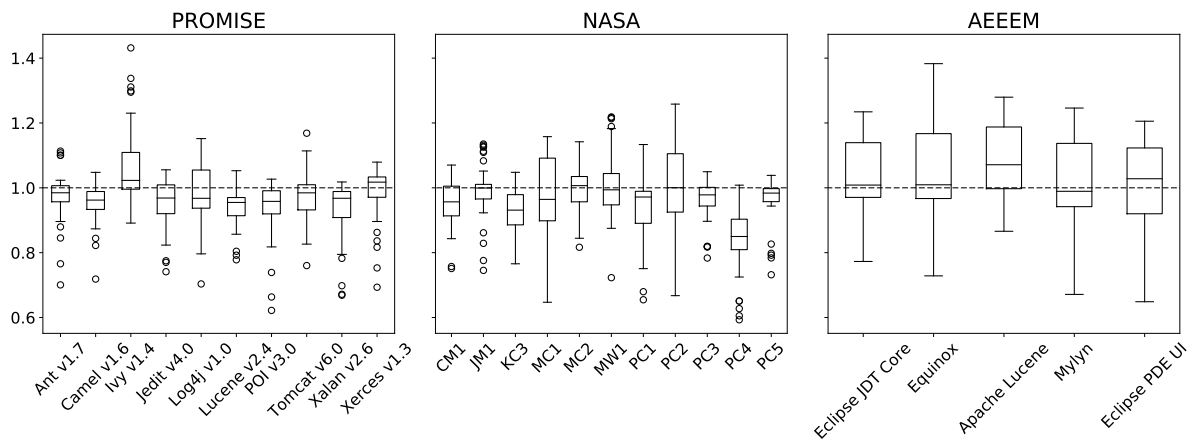
**The supervised models with feature reduction techniques significantly outperform the unsupervised models with feature reduction techniques.** Figure 5.2 shows the AUC after applying the feature reduction techniques to the supervised and unsupervised models. The supervised models significantly outperform the unsupervised models. Prior research [41] reported that spectral clustering (SC) is the only studied unsupervised defect prediction model that outperforms the supervised models.

The reason that the unsupervised models perform worse than the supervised models in Figure 5.2 is that we consider all the unsupervised models together, to be able to provide a more generic conclusion. However, as Figure 5.2 shows, some unsupervised defect prediction models perform better than others.

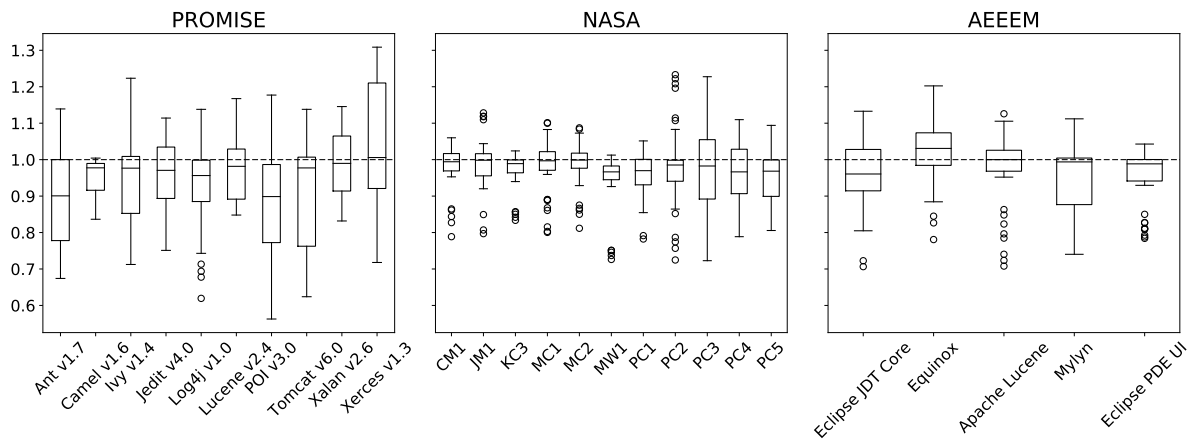
**In the AEEEM dataset, the feature reduction techniques improve the prediction performance of the supervised models for most projects.** We observe that the feature reduction techniques did not improve the prediction performance in many projects, as the median values of several boxplots in Figure 5.3 are lower than 1.0. However, the studied feature reduction techniques improved the prediction performance of the supervised models for many projects in



**Figure 5.2 The Scott-Knott ESD test results for both the supervised and unsupervised models. We use two colors (dark gray and light gray) and two lines (solid and dashed lines) to represent the ranks according to the Scott-Knott ESD test. The adjacent boxplots with the same color and line indicate the same rank. Otherwise, the rank is changed at that point. The different rank indicates a statistical significant difference with small effect size or over. Each boxplot has 130 median AUC values (5 defect prediction models times 26 projects). The x-axis refers to the feature reduction techniques; the y-axis refers to the AUC values. In the x-axis, the “SVL\_”-prefix refers to the 5 supervised defect prediction models; the “USVL\_”-prefix refers to the 5 unsupervised defect prediction models.**



(a) The supervised models



(b) The unsupervised models

**Figure 5.3** The ratios of the AUCs of the supervised and unsupervised prediction models. Each boxplot contains 40 ratio values (5 prediction models \* 8 feature reduction techniques). The x-axis indicates the studied projects. The y-axis indicates the ratio. The dashed line indicates a ratio of 1.0. A ratio larger than 1.0 indicates that the feature reduction technique improved the AUC compared to the original models, while a ratio smaller than 1.0 indicates that the feature reduction technique reduced the AUC.

the AEEEM dataset (Figure 5.3(a)). We further investigate this phenomenon in Section 5.3.1.

**The neural network-based techniques improve the prediction performance of the supervised/unsupervised prediction models for most projects.** Table 5.1 and Table 5.2 show the median ratios for the feature reduction techniques. We observe that the neural network-based feature reduction techniques RBM and AE have the most gray cells for the supervised/unsupervised prediction models in combination with the feature reduction techniques.

However, almost all feature reduction techniques did not improve the prediction performance in the NASA dataset except for FA with the unsupervised prediction models. FA combined with the unsupervised prediction models improved over half of the projects in the NASA dataset. We further investigate why the feature reduction techniques work well for the AEEEM dataset but not for the other datasets in Section 5.3.1.

## **5.2 RQ2: What is the impact of feature reduction techniques on the variance of the performance across defect prediction models?**

*Motivation:* A challenge in applying defect prediction for practitioners is to select the best-performing model for their data from many possible defect prediction models [11, 77]. In this RQ, we studied the variance in performance across defect prediction models of the studied feature reduction techniques for a particular dataset. If this variance is small, the practitioner does not need to worry about the choice at all, as the models perform similarly across datasets.

*Approach:* We used the interquartile range (IQR) which captures the performance variance across the studied defect prediction models for a given project and feature reduction technique. We used the AUC values of all the studied supervised and unsupervised models for all bootstrap samples to conduct a new *bootstrap sampling* to calculate the IQR for each reduction technique and each project. We calculated an IQR value as follows:

1. Sample 100 AUC values at random from the 100 AUC values for each studied supervised/unsupervised model allowing for replacement.
2. Compute the median AUC value across the sampled 100 AUC values.
3. Repeat steps 1 and 2 100 times.

**Table 5.1 The median AUC ratios of the feature reduction and selection techniques in the studied supervised models. A ratio larger than 1 indicates that the feature reduction/selection technique improved the AUC compared to the original models. The gray cells refer to the ratios that are greater than 1.0. The “Improved” row indicates the number of projects for which a feature reduction/selection technique improved the performance.**

		RBM	AE	PCA	FM	FA	RP	TCA	TCA+	CFS	ConFS
PROMISE	Ant v1.7	1.015	1.004	0.983	0.896	0.995	0.956	0.962	0.962	1.012	1.006
	Camel v1.6	0.973	0.979	0.921	0.880	0.996	0.962	0.948	0.948	0.962	0.956
	Ivy v1.4	1.065	1.081	1.005	1.005	1.026	0.999	1.036	1.020	1.062	1.037
	Jedit v4.0	1.012	0.996	0.971	0.775	1.011	0.921	0.952	0.947	0.984	0.996
	Log4j v1.0	1.047	1.034	1.000	0.960	1.012	0.895	0.961	0.954	0.993	1.002
	Lucene v2.4	1.016	0.956	0.940	0.804	0.978	0.886	0.962	0.955	0.990	0.978
	POI v3.0	0.932	0.899	0.955	0.739	0.993	0.946	0.965	0.971	1.016	1.003
	Tomcat v6.0	1.014	1.016	0.947	0.875	0.997	0.925	0.985	0.990	1.045	1.026
	Xalan v2.6	0.861	0.912	0.987	0.698	0.992	0.970	0.993	0.985	1.012	0.998
	Xerces v1.3	1.032	1.022	0.969	0.817	1.019	1.017	1.034	1.033	1.005	1.012
NASA	CM1	1.004	0.979	0.939	0.984	0.974	0.995	0.949	0.940	1.028	0.981
	JM1	1.004	1.001	0.997	0.923	0.999	1.003	1.008	0.956	1.000	1.001
	KC3	0.910	0.944	0.901	0.923	1.010	0.924	0.905	0.874	1.025	0.985
	MC1	0.956	1.021	0.932	0.901	0.999	0.915	0.980	0.932	1.001	0.973
	MC2	1.019	1.014	0.955	0.947	1.037	1.014	0.976	0.972	0.973	0.968
	MW1	1.005	1.016	0.949	0.984	1.019	0.991	0.952	0.963	1.016	1.015
	PC1	0.906	0.830	0.936	0.972	1.013	0.976	0.971	0.925	1.004	1.021
	PC2	1.039	1.019	0.931	0.994	1.020	0.999	0.935	0.933	1.169	1.028
	PC3	0.944	0.971	0.985	0.921	0.997	0.988	1.001	0.968	1.026	0.995
	PC4	0.793	0.806	0.931	0.651	0.906	0.862	0.866	0.870	1.003	0.986
PC5	0.999	0.982	0.994	0.798	0.990	0.977	0.989	0.986	0.986	0.997	
AIEEE	Eclipse JDT Core	1.028	1.010	1.019	0.874	1.007	0.959	1.015	0.986	0.993	1.012
	Equinox	1.074	1.043	1.006	0.973	1.028	0.996	1.000	1.010	1.060	1.024
	Apache Lucene	1.112	1.079	1.090	1.045	1.009	1.063	1.030	1.039	1.021	1.035
	Mylyn	1.089	1.057	0.980	0.921	1.052	0.951	1.066	1.065	1.037	1.023
	Eclipse PDE UI	1.071	1.058	1.075	0.861	1.035	0.944	1.020	0.905	0.995	1.006
Improved	17	15	5	2	14	4	8	5	17	15	

**Table 5.2 The median AUC ratios of the feature reduction and selection techniques in the studied unsupervised models.**

	RBM	AE	PCA	FM	FA	RP	TCA	TCA+	CFS	ConFS	
PROMISE	Ant v1.7	1.005	1.000	1.001	0.739	0.896	0.695	0.918	0.748	1.002	1.006
	Camel v1.6	0.989	0.980	1.002	0.894	0.979	0.848	0.984	0.933	0.977	0.988
	Ivy v1.4	1.040	1.011	1.006	0.938	0.850	0.750	0.991	0.874	0.937	0.928
	Jedit v4.0	1.055	1.038	1.000	0.895	1.033	0.888	0.966	0.802	0.939	0.962
	Log4j v1.0	0.993	0.989	1.001	0.950	0.973	0.694	0.924	0.790	0.954	0.986
	Lucene v2.4	1.110	1.131	1.001	0.870	0.989	0.894	0.966	0.858	0.984	1.004
	POI v3.0	0.919	0.924	1.002	0.755	1.036	0.875	0.610	0.755	0.962	0.982
	Tomcat v6.0	1.041	1.007	0.997	0.762	0.956	0.652	1.007	0.763	1.010	0.968
	Xalan v2.6	1.070	1.065	0.997	0.885	0.936	0.965	1.084	0.848	0.964	0.974
	Xerces v1.3	1.244	1.210	1.008	0.885	1.259	0.893	1.201	0.939	1.106	1.017
NASA	CM1	1.008	0.996	0.997	0.967	0.979	1.049	0.983	0.844	1.030	0.949
	JM1	0.992	1.023	1.002	0.953	1.036	0.997	1.011	0.849	1.008	1.001
	KC3	0.970	0.987	1.000	0.974	1.014	0.997	0.998	0.850	0.968	0.955
	MC1	1.005	1.023	1.000	0.891	1.021	1.000	0.994	0.816	1.081	1.023
	MC2	1.073	1.043	1.000	0.981	1.018	0.985	0.995	0.862	0.993	0.991
	MW1	0.939	0.971	1.000	0.973	0.966	0.990	0.957	0.746	1.017	1.019
	PC1	0.990	0.996	0.996	0.927	1.027	0.935	0.993	0.855	1.102	1.041
	PC2	0.985	0.976	0.994	0.975	0.985	0.993	0.966	0.774	1.035	0.940
	PC3	1.074	1.090	1.000	0.940	0.954	0.758	1.041	0.831	1.213	1.123
	PC4	0.986	0.982	0.998	0.902	1.062	0.829	0.961	0.818	1.123	0.991
PC5	0.971	0.975	0.999	0.866	0.971	0.987	0.932	0.857	1.030	0.997	
AEEEM	Eclipse JDT Core	1.121	1.089	0.997	0.888	0.957	0.995	1.107	0.805	1.017	1.007
	Equinox	1.029	1.105	1.000	0.913	1.069	1.033	1.073	0.900	1.025	1.033
	Apache Lucene	1.031	1.053	1.000	0.863	0.989	1.006	1.024	0.740	0.956	0.981
	Mylyn	1.006	1.015	1.001	0.829	0.996	0.892	0.994	0.827	1.009	0.996
	Eclipse PDE UI	1.015	1.024	1.000	0.827	0.980	0.977	0.993	0.793	1.003	1.024
Improved	16	15	9	0	10	3	8	0	16	11	



4. Compute the IQR value for the 500 sampled median AUC values (5 supervised/unsupervised prediction models \* 100 median AUC values) for each feature reduction/selection technique for each studied project.

Where the IQR values are computed as follows:

$$\text{IQR} = Q_3 - Q_1$$

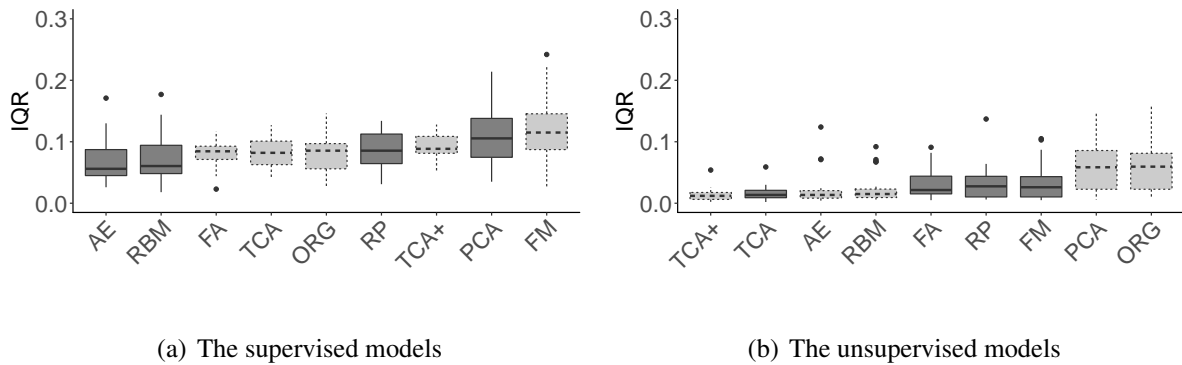
where  $Q_1$  is the first quartile of the 500 sampled median AUC values, and  $Q_3$  is the third quartile of the 500 sampled median AUC values. The first quartile is the median between the smallest and the median of the 500 AUC values, and the third quartile is the median between the median and the largest of the 500 AUC values. As we studied 26 projects, we have 26 IQR values for each feature reduction technique. We used the Scott-Knott ESD test to compare the distributions of IQRs for each feature reduction technique. Figure 5.4 shows the results of the Scott-Knott ESD test.

In addition, we compared the IQR values of the prediction models across the feature reduction techniques for each project. Table 5.3 and Table 5.4 show the results of the IQR analysis at the project level. Each cell contains an IQR value that was computed from 500 bootstrapped median AUC values of the supervised and unsupervised prediction models.

**Results:** **The neural network-based feature reduction techniques, RBM and AE, generate features that result in less variance across the supervised models than the original features.** Figure 5.4(a) shows that the original features (ORG) are in the second rank, and RBM and AE belong to the first rank. Hence, RBM and AE significantly improve the variance of the performance across the supervised defect prediction models.

**Almost all feature reduction techniques (except PCA) generate features that have a significantly smaller performance variance across the unsupervised models than the original features.** Figure 5.4(b) shows that the unsupervised models that use features that were generated by PCA, or the original features are in the lowest rank. Hence, using the feature reduction techniques (except PCA) in combination with an unsupervised defect prediction model results in a small performance variance, which is helpful for practitioners.

**The neural network-based feature reduction techniques improve the performance variance of the supervised models for the largest number of projects.** Table 5.3 shows that the



**Figure 5.4 The Scott-Knott ESD test results for the IQR of the supervised and unsupervised models. We use two colors (dark gray and light gray) and two lines (solid and dashed lines) to represent the ranks according to the Scott-Knott ESD test. The adjacent boxplots with the same color and line indicate the same rank. Otherwise, the rank is changed at that point. The different rank indicates a statistical significant difference with small effect size or over in variance (IQR). Each boxplot has 26 IQR values (one for each project). The x-axis refers to the feature reduction techniques; the y-axis refers to the IQR values.**

**Table 5.3 The IQR ratio for the feature reduction techniques in the studied supervised prediction models. The gray cells refer to the ratios that are greater than 1.0. The “Improved” row indicates the number of gray cells in the column.**

		RBM	AE	PCA	FM	FA	RP	TCA	TCA+	CFS	ConFS
PROMISE	Ant v1.7	1.769	1.330	0.849	0.439	1.024	0.943	0.975	1.006	1.294	1.031
	Camel v1.6	3.885	3.027	0.842	0.809	0.774	0.777	1.312	1.357	1.637	1.466
	Ivy v1.4	1.841	2.964	1.978	2.197	2.001	1.151	3.641	1.233	1.340	2.072
	Jedit v4.0	1.018	1.180	0.562	0.501	0.771	0.595	0.688	0.718	0.767	0.683
	Log4j v1.0	1.335	2.643	1.030	5.212	1.388	1.898	1.252	1.189	1.361	1.377
	Lucene v2.4	1.737	0.661	0.554	0.473	0.635	1.141	0.666	0.647	0.800	0.658
	POI v3.0	0.973	0.586	0.628	0.241	0.959	1.642	0.894	0.977	1.268	1.157
	Tomcat v6.0	1.351	1.429	0.671	0.657	1.096	0.739	0.854	0.921	1.369	1.204
	Xalan v2.6	0.762	0.615	0.863	0.495	0.725	0.750	0.531	0.556	0.972	0.841
	Xerces v1.3	1.918	1.749	0.806	1.022	0.847	0.943	1.064	1.078	1.171	1.417
NASA	CM1	0.196	0.227	0.241	0.311	0.329	0.683	0.389	0.430	0.589	0.453
	JM1	2.502	2.579	1.886	0.604	0.804	0.972	1.036	0.758	0.892	0.986
	KC3	0.582	0.955	1.509	1.130	0.661	1.093	0.906	0.751	1.871	1.710
	MC1	0.639	0.670	0.605	0.734	0.980	0.810	1.227	0.967	0.831	2.011
	MC2	1.087	1.080	0.545	1.074	0.808	0.835	0.734	0.679	1.121	1.169
	MW1	1.199	1.451	0.722	4.598	2.933	1.166	2.113	1.697	1.792	1.274
	PC1	0.826	0.984	0.649	1.078	1.326	1.936	2.501	0.875	1.179	1.578
	PC2	0.541	0.565	0.454	0.788	1.067	1.317	0.937	0.898	0.796	0.389
	PC3	1.060	1.522	0.684	0.985	1.274	0.718	0.862	0.695	1.331	1.145
	PC4	0.885	1.504	0.548	0.397	0.852	1.055	0.681	0.750	5.539	1.138
PC5	0.796	1.892	0.752	0.344	0.908	0.705	0.723	0.757	0.798	1.009	
AEEEM	Eclipse JDT Core	1.332	1.433	1.246	0.434	1.556	0.998	1.373	0.970	1.072	1.393
	Equinox	1.989	2.927	1.543	0.821	1.076	0.822	1.075	1.129	1.438	1.337
	Apache Lucene	1.373	1.844	0.936	0.608	0.989	1.678	1.014	0.993	0.881	1.022
	Mylyn	5.981	4.024	0.940	0.946	5.923	1.155	1.289	1.282	1.992	2.406
	Eclipse PDE UI	1.574	1.208	1.096	0.224	0.392	0.374	0.434	0.384	0.475	0.452
Improved	17	18	7	7	11	11	12	8	16	19	

**Table 5.4 The IQR ratio for the feature reduction techniques in the studied unsupervised prediction models. The gray cells refer to the ratios that are greater than 1.0. The “Improved” row indicates the number of gray cells in the column.**

	RBM	AE	PCA	FM	FA	RP	TCA	TCA+	CFS	ConFS	
PROMISE	Ant v1.7	2.693	1.772	1.035	2.539	3.228	3.217	1.447	4.320	0.335	0.245
	Camel v1.6	1.224	1.715	1.331	0.961	1.659	0.778	0.503	0.680	0.209	0.301
	Ivy v1.4	2.881	4.073	0.924	2.096	3.690	3.628	3.317	3.076	2.809	2.153
	Jedit v4.0	3.221	3.229	1.012	1.704	4.387	2.213	3.346	4.496	0.671	0.790
	Log4j v1.0	3.400	3.445	1.221	1.346	0.542	0.977	1.562	0.783	0.762	0.793
	Lucene v2.4	0.460	0.982	0.653	0.796	0.577	1.865	0.945	0.560	0.176	0.163
	POI v3.0	0.764	0.582	1.152	7.759	1.425	7.589	4.332	2.842	0.415	0.927
	Tomcat v6.0	2.993	1.985	1.116	2.326	3.566	4.079	1.512	5.814	0.630	0.904
	Xalan v2.6	2.028	3.317	0.763	2.339	1.363	3.747	0.829	0.907	0.362	2.860
	Xerces v1.3	8.508	6.308	1.132	11.297	6.729	10.673	2.089	10.157	4.720	0.980
NASA	CM1	2.032	2.820	1.074	0.694	1.274	1.535	1.573	3.922	1.238	0.688
	JM1	1.206	5.044	0.988	0.983	1.360	1.304	39.761	154.852	1.210	1.009
	KC3	0.920	0.753	0.936	1.280	0.639	0.780	0.946	3.250	0.591	0.750
	MC1	5.284	3.674	1.035	0.576	3.833	1.261	5.611	3.635	1.132	0.971
	MC2	1.377	2.056	1.839	0.778	1.670	1.919	2.709	4.491	1.457	1.441
	MW1	2.059	1.207	0.952	1.043	0.897	1.008	1.325	1.520	0.814	0.816
	PC1	2.413	2.298	0.996	1.047	1.405	0.866	3.256	3.491	0.964	0.748
	PC2	2.058	1.900	1.014	2.441	5.610	1.035	5.055	7.867	2.231	1.436
	PC3	2.293	2.230	1.069	2.318	4.411	7.025	13.146	13.203	2.622	1.182
	PC4	4.032	3.431	0.917	4.188	1.266	3.476	9.617	10.161	1.653	1.177
PC5	5.900	11.856	1.025	18.438	4.839	2.096	30.503	5.129	1.355	1.267	
AEEEM	Eclipse JDT Core	15.497	15.262	0.973	1.727	1.290	3.088	20.680	8.202	1.104	1.505
	Equinox	3.504	7.741	0.684	3.573	1.155	1.420	6.460	9.422	0.676	0.601
	Apache Lucene	8.221	6.507	0.989	0.521	4.423	1.610	5.717	8.040	0.659	0.920
	Mylyn	10.940	9.597	1.038	8.177	1.099	1.036	17.600	11.010	0.980	0.889
	Eclipse PDE UI	3.237	2.806	1.068	0.690	1.235	0.465	5.627	6.956	0.389	0.366
Improved	23	23	15	18	22	21	22	22	11	9	

features that were generated by the neural network-based feature reduction techniques (RBM and AE) improved the performance variance (IQR) across the studied supervised prediction models for the largest number of projects compared to the other feature reduction techniques, and the original models. RBM and AE also belong to the first rank of the overall performance variance result (Figure 5.4(a)).

**The neural network-based feature reduction techniques improve the performance variance of the unsupervised models for the largest number of projects.** Table 5.4 shows that the features that were generated by the neural network-based feature reduction techniques (RBM and AE) improved the performance variance across the studied unsupervised prediction models for the largest number of projects. Interestingly, in terms of overall performance variance, TCA and TCA+ belong to the first and the second rank (Figure 5.4(b)). However, the difference with the neural network-based feature reduction techniques is only small (TCA+) and negligible (TCA), according to the Cliff's delta effect size.

### **5.3 RQ3: How do feature selection techniques compare to feature reduction techniques when applied to defect prediction?**

In this RQ, we compare feature reduction and selection techniques along two dimensions: the performance and the performance variance of the defect prediction models. We study the correlation-based (CFS) and consistency-based (ConFS) feature selection techniques, as they performed best according to prior studies [9, 18].

*Motivation:* In RQ1 and RQ2, we found that several feature reduction techniques (FA, RBM and AE) outperform the original features (ORG) in terms of performance or performance variance of the defect prediction models. Prior work [9, 18] showed that several feature selection techniques outperform the original models as well. In this RQ, we compare the performance (AUC) and the performance variance (IQR) of the feature reduction and selection techniques when applied to defect prediction models.

*Approach:* The experimental procedure is the same as the procedures of RQ1 and RQ2 (only we use the two feature selection techniques CFS and ConFS instead of the feature reduction techniques).

*Results:* **The feature selection techniques (correlation-based feature selection (CFS) and consistency-based feature selection technique (ConFS)) significantly outperform the original features (ORG) in the supervised models, and perform as well as the feature agglomeration (FA) reduction technique.** Figure 5.5(a) shows the AUC values and the results of the Scott-Knott ESD test for the supervised models after applying the studied feature reduction and selection techniques. <sup>(\*2)</sup> Each boxplot shows the median AUC values for the projects using a certain feature reduction/selection technique. CFS, ConFS and FA are in the highest rank by themselves, which indicates that the subsets of features that were selected by CFS or ConFS perform as well as the feature sets that were generated by FA for the supervised models.

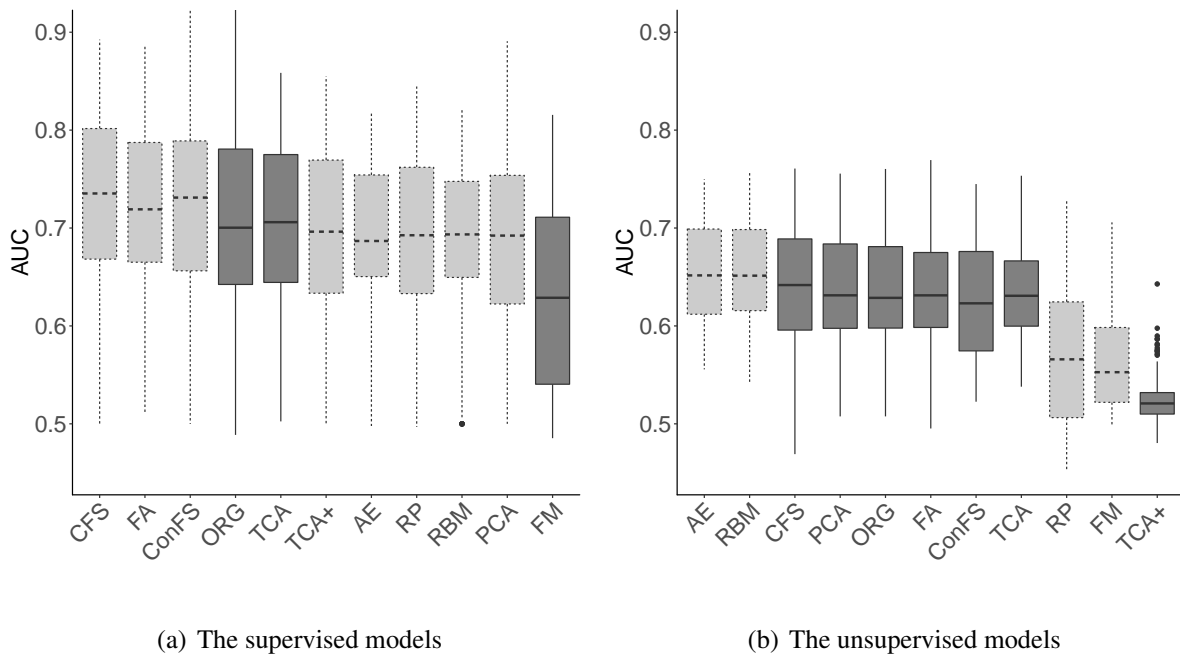
**The neural network-based feature reduction techniques (RBM and AE) significantly outperform the feature selection techniques (CFS and ConFS) for the unsupervised defect prediction models.** The highest rank contains only the two studied neural network-based feature reduction techniques (Figure 5.5(b)). The studied feature selection techniques (CFS and ConFS) belong to the second rank, together with the original models (ORG). Hence, the studied feature selection techniques have a worse performance than the neural network-based feature reduction techniques for the unsupervised defect prediction models.

**In the supervised models, applying the neural network-based feature reduction techniques, RBM and AE, or the feature selection techniques, CFS and ConFS, significantly outperforms the original models in terms of performance variance.** The original models (ORG) belong to the third rank (Figure 5.6(a)). The neural network-based feature reduction techniques and the feature selection techniques belong to the first or second rank, hence they have a smaller performance variance than the original models.

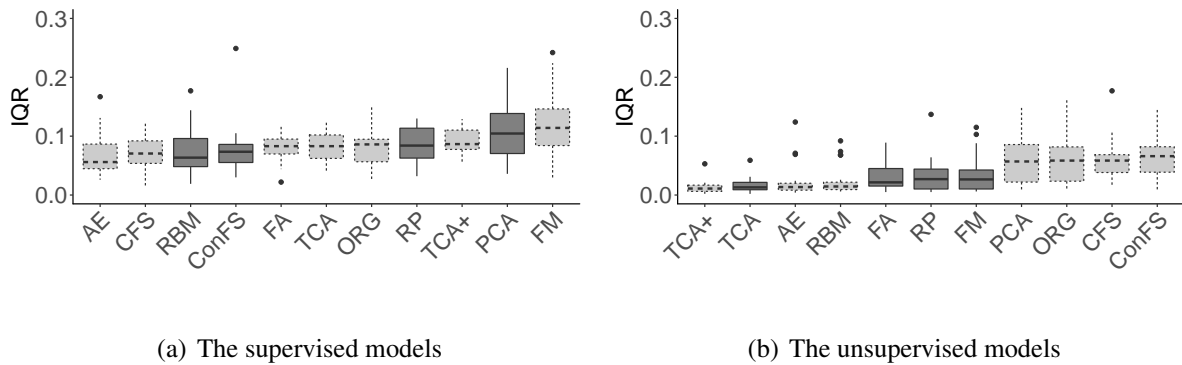
**In the unsupervised models, all feature reduction techniques (except PCA) significantly outperform the feature selection techniques in terms of performance variance.** The feature selection techniques belong to the worst rank together with the original models (ORG) and PCA (Figure 5.6(b)). Hence, the studied feature selection had a larger performance variance than almost all the studied feature reduction techniques for the unsupervised defect prediction

---

(\*2): Note that the ranks are slightly different from Figure 5.1 due to the fact that Scott-Knott ESD is a clustering algorithm, and hence affected by the total set of input distributions. For more information see <https://github.com/klainfo/ScottKnottESD>.



**Figure 5.5 The Scott-Knott ESD test results for the supervised (logistic regression, random forest, naive Bayes, decision tree, and logistic model tree) and unsupervised (spectral clustering, k-means, partition around medoids, fuzzy C-means, neural-gas) models. We use two colors (dark gray and light gray) and two lines (solid and dashed lines) to represent the ranks according to the Scott-Knott ESD test. The adjacent boxplots with the same color and line indicate the same rank. Otherwise, the rank is changed at that point. The different rank indicates a statistical significant difference with small effect size or over in performance. Each boxplot has 130 median AUC values (5 defect prediction models times 26 projects). The x-axis refers to the feature reduction/selection techniques; the y-axis refers to the AUC values.**



**Figure 5.6 The Scott-Knott ESD test results for IQR in the supervised and unsupervised models. We use two colors (dark gray and light gray) and two lines (solid and dashed lines) to represent the ranks according to the Scott-Knott ESD test. The adjacent boxplots with the same color and line indicate the same rank. Otherwise, the rank is changed at that point. The different rank indicates a statistical significant difference with small effect size or over in variance (IQR). Each boxplot has 26 IQR values (one for each project). The x-axis refers to the feature reduction/selection techniques; the y-axis refers to the IQR values.**



models.

**Our above findings for RQ3 are confirmed by our project-level analysis.** Table 5.1 and Table 5.2 show the median ratios of the performance of each feature reduction/selection technique compared to the original models. We calculated this ratio as follows:

$$\text{ratio} = \frac{AUC_{FRS}}{AUC_{ORG}}$$

Where  $AUC_{ORG}$  is the AUC (the median AUC across all bootstrap samples) of a prediction model using the original features, and  $AUC_{FRS}$  is the AUC of a prediction model using the features that were generated/selected by a particular feature reduction or selection technique. We computed the median ratio across the five studied supervised and unsupervised prediction models. Table 5.1 and Table 5.2 confirm our above findings about the performance of the studied feature selection techniques compared to that of the feature reduction techniques.

Table 5.3 and Table 5.4 show the IQR ratio values of each feature reduction/selection technique. We define this ratio as follows:

$$\text{ratio} = \frac{IQR_{FRS}}{IQR_{ORG}}$$

Where  $IQR_{ORG}$  is the IQR (the median IQR across all bootstrap samples) of a prediction model using the original features, and  $IQR_{FRS}$  is the IQR of a prediction model using the features that were generated/selected by a particular feature reduction or selection technique. We calculated the median IQR value for the supervised models using bootstrap samples as follows:

1. Sample 100 values following the distribution of the 100 AUC values for each studied supervised model while allowing for replacement.
2. Compute the median AUC value across the sampled 100 values.
3. Repeat steps 1 and 2 100 times.
4. Compute the IQR value for the 500 sampled median AUC values (5 supervised prediction models \* 100 median AUC values) for each feature reduction/selection technique for each studied project.

We repeated the above procedure for the unsupervised models. Table 5.3 shows that the project-level results confirm our findings above, as the RBM and AE feature reduction techniques and the CFS and ConFS feature selection techniques improve the performance variance

of most projects compared to the other techniques. In addition, Table 5.4 shows that all feature reduction techniques improve the performance variance of more projects than the CFS and ConFS feature selection techniques.

### 5.3.1 Why do feature reduction techniques work well in the AEEEM dataset?

*Motivation:* We observed that the feature reduction techniques work better for the projects in the AEEEM dataset than for the projects in the other datasets. Ghotra et al. [9] applied PCA to the data of each project to capture its richness. We use the same analysis to investigate whether the dataset richness is an explanation of why feature reduction techniques work better for the AEEEM dataset.

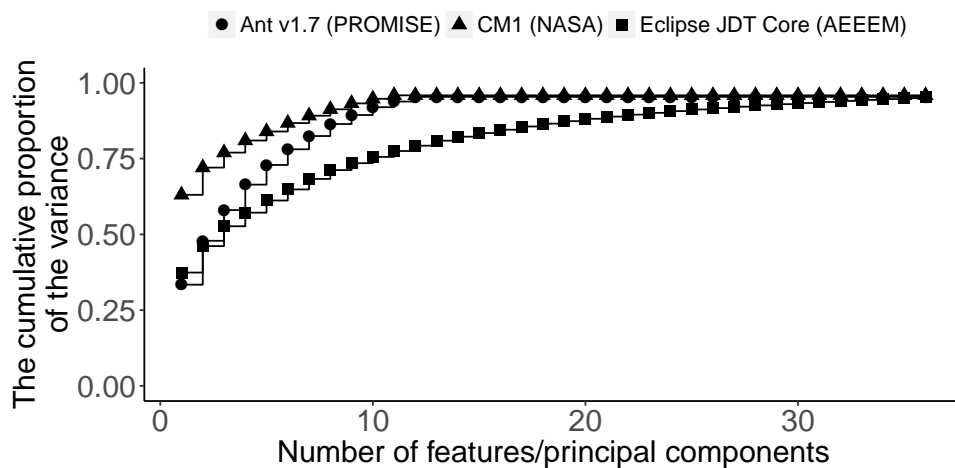
*Approach:* The idea behind Ghotra et al.'s analysis [9] is to generate features from a dataset using PCA that (together) retain at least 95% of the variance of the original dataset. Ghotra et al. reason that a larger number of generated features indicates a richer dataset. Likewise, they interpret that a small number of generated features indicates redundancy in the original dataset. We applied PCA to each project and counted the number of generated features.

*Results:* **The PROMISE, NASA, and AEEEM datasets have different data richness characteristics, however; the characteristics of the projects within each dataset are consistent.** Table 5.5 shows the number of generated features. While the number of generated features is approximately the same for the PROMISE and NASA projects, the proportion of generated features compared to the number of original features is different. In addition, this proportion is even lower for the AEEEM projects. Hence, we conclude that the datasets have different characteristics in terms of data richness. However, within each dataset, the projects have approximately the same richness characteristics.

**The original features of the projects in the AEEEM dataset are more diverse than the projects of the other datasets.** We observe that 36 principal components are needed to cover 95% of the variance of the Eclipse JDT Core project in Figure 5.7, compared to 12 components for the Ant project and 11 for the CM1 project. Hence, the original features of the AEEEM dataset are much more diverse than those of the PROMISE and the NASA datasets. The diversity of the AEEEM dataset could be a reason why feature reduction techniques improve the performance of this dataset.

**Table 5.5 The number of generated features (principal components) that are needed to account for 95% of the data variance.**

<b>Studied Dataset</b>	<b>Studied Project</b>	<b># of Studied Features</b>	<b># of Generated Features</b>	<b>% of Generated Features</b>
PROMISE	Ant v1.7	20	12	60.0
	Camel v1.6	20	12	60.0
	Ivy v1.4	20	10	50.0
	Jedit v4.0	20	12	60.0
	Log4j v1.0	20	12	60.0
	Lucene v2.4	20	12	60.0
	POI v3.0	20	12	60.0
	Tomcat v6.0	20	12	60.0
	Xalan v2.6	20	12	60.0
	Xerces v1.3	20	12	60.0
NASA	CM1	37	11	29.7
	JM1	21	8	38.1
	KC3	39	10	25.6
	MC1	38	15	39.5
	MC2	39	11	28.2
	MW1	37	11	29.7
	PC1	37	12	32.4
	PC2	36	10	27.8
	PC3	37	13	35.1
	PC4	37	14	37.8
	PC5	38	15	39.5
AEEEM	Eclipse JDT Core	212	36	17.0
	Equinox	212	31	14.6
	Apache Lucene	212	33	15.6
	Mylyn	212	46	21.7
	Eclipse PDE UI	212	38	17.9



**Figure 5.7** The number of principal components (features that were generated by PCA) that are needed to account for the original data variance for the Ant, CM1 and Eclipse projects in the PROMISE, NASA and AEEEM datasets. The x-axis indicates the number of principal components. The y-axis indicates the cumulative proportion of the variance. The other projects of the datasets showed a similar pattern.

### 5.3.2 Comparing feature selection and reduction techniques along the dimensions of understandability and execution time

The understandability of the metrics in a defect prediction model is important, as understandable metrics make the model, and its predictions, easier to explain [78]. Feature reduction techniques combine the original features into one or more newly-generated features. Hence, these newly-generated features are by definition harder to understand than the features that are a subset (i.e., they were *selected*) of the original features. We inspected the feature sets that were generated during our experiments, and we observed that almost all generated feature sets consist of features that are a complex combination of all available original features. Hence, defect prediction models that are generated using feature reduction techniques are harder to understand than those that use feature selection.

In addition, the execution time of a feature reduction/selection technique and a defect prediction model is important – models that take too long to build or execute are not very useful in practice. To conduct our experiments in a timely manner, we ran them on a cluster of servers in parallel. Hence, it is difficult to compare the execution time of the experiments. In general, the execution time of our feature reduction/selection techniques and defect prediction models was short (i.e., in the range of minutes). Therefore, execution time is not a very problematic metric for practitioners who wish to apply feature reduction or selection to their defect prediction models.

## 6. Discussion: Which features are generated by the feature reduction techniques?

In RQ1, RQ2, and RQ3, we observed that some feature reduction/selection techniques generate/select features that perform defect prediction better and less variance than the features that were generated/selected by other feature reduction/selection techniques. In particular, we found that RBM and AE outperform the other studied feature reduction/selection techniques for the unsupervised models. However, RBM and AE are less-performing feature reduction techniques than the original models (ORG) for the supervised models. In this discussion, we take a closer look at the generated features to investigate why neural network-based feature reduction techniques perform well for the unsupervised defect prediction models, but not for the supervised models. In this section, we discuss possible explanations for the differences in AUC and variance of the performance.

*Approach:* We focused our discussion on the RBM and AE feature reduction techniques, as these techniques generate new features by assigning (combinations of) weights to the original features. For example, a newly generated feature may be generated by 0.5 times original metric 1 and 0.5 times original metric 2. These weight sets allow us to study how the new features are related to the original features, and to the features that were generated by other feature reduction techniques, and extract possible explanations for the improved and small variance of the performance. RP and PCA also generate new features by assigning weights to the original features, however, the weights of RP are randomly generated, and PCA generates a different number of features for each project, which makes them difficult to compare. Hence, we focused our discussion on RBM and AE. As we generated 100 new feature sets of 10 features using RBM and AE, we generated 1,000 weight sets using these two feature reduction techniques for each project. For each feature reduction technique, we randomly selected 10 (out of the 100) generated feature sets for our investigation.

We conducted correlation analysis and clustering analysis on the studied feature sets to study their similarity within and across projects. The correlation analysis shows how independent the features that are generated for a project are. Highly correlated features can negatively impact


the performance of regression models [7] and this effect can affect our supervised models as well. We first calculated the Spearman rank correlation [79] between the generated features in a feature set within a project. Each generated feature was normalized using the  $z$ -score. We chose Spearman rank correlation because it is non-parametric, and therefore requires no assumption about the distribution of the studied data.

To study the similarity of the generated feature sets across projects within a dataset, we compared the weight sets of the generated features. First, we normalized all weights using  $z$ -score normalization. Second, we used  $k$ -means to cluster the weight sets of the features, and then we used t-distributed stochastic neighbour embedding (t-SNE) [80] to visualize the clustering results. t-SNE is commonly used for visualizing high dimensional features in scatter plots [81]. In particular, t-SNE models high-dimensional objects (i.e., feature sets) by two- or three-dimensional points such that similar objects are close, and dissimilar objects are further away from each other.

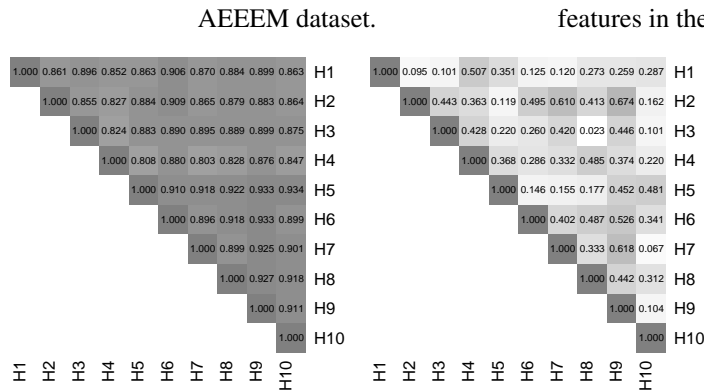
The goal of our clustering analysis is to find out how similar the generated features are across projects within a dataset. Hence, we configured  $k$ -means to search for 10 clusters (as we are generating 10 new features). We visualized the clustering results using the default settings of t-SNE.

*Results:* **RBM generates feature sets in which all features are strongly correlated with each other.** Figure 6.1(a) shows the Spearman rank correlation of one set of RBM-generated features for the Ant project in the PROMISE dataset using a heatmap. We observe that all correlations are close to 1 (dark gray), which means that all RBM-generated features in the feature set are strongly correlated (and hence similar) to each other. We observe similar correlations for the other studied feature sets for the PROMISE and NASA datasets. RBM generates weakly correlated features for several projects in the AEEEM dataset (e.g., for the Eclipse project: Figure 6.1(b)). However, if we use a smaller set of original features from that dataset, such as only the change features, RBM generates strongly correlated features for these projects as well (e.g., for the Eclipse project: Figure 6.1(c)).

We observe similar correlations in feature sets that were generated by AE. However, the correlation within the AE-generated feature sets appeared to be linked to the specific project. For example, AE generates sets of features that are strongly correlated to each other for the



(a) RBM-generated features in the Ant project in the PROMISE dataset. (b) RBM-generated features in the Eclipse project in the PROMISE dataset. (c) RBM-generated features in the Eclipse project using change dataset.



(d) AE-generated features in the Log4j project in the PROMISE dataset. (e) AE-generated features in the Camel project in the PROMISE dataset.

**Figure 6.1 The Spearman rank correlation of generated features in the studied datasets. The darker colours indicate a strong absolute correlation (close or equal to one). The lighter colours indicate a weak absolute correlation (close or equal to zero).**




Log4j project (see Figure 6.1(d)), but features that are not as strongly correlated for the Camel project (see Figure 6.1(e)). Hence, a possible explanation of the reason for the small variance of the performance of features that were generated by the neural network-based feature reduction techniques (i.e., RBM and AE) could be the strong correlation within the generated feature sets.

**RBM mostly generates the same feature sets across projects.** Figure 6.2(a) shows the  $k$ -means clustering result for one bootstrap sample of the RBM-generated weight sets in the PROMISE dataset using t-SNE [81]. Each shape refers to a project, and each color refers to a cluster that was identified by  $k$ -means clustering. The x-axis and y-axis refer to the t-SNE features that were generated from the 260 RBM-generated weight sets (10 weight sets on each project) by t-SNE. Hence, if there are 10 clearly identifiable groups of different shapes with the same colour in the t-SNE plot, we can conclude that the generated weight sets (and hence the features) are the same across projects. We observe that each colored cluster contains all shapes, which indicates that each cluster contains all projects. Hence, in this bootstrap sample, RBM generated the same feature sets across projects in the PROMISE dataset. We noticed a similar pattern for the other projects in the PROMISE dataset for the other bootstrap samples.


Figure 6.2(b) shows the result for the RBM-generated weight sets in the NASA dataset. The figure shows that RBM generated different feature sets across projects in the NASA dataset. A possible reason is that the original features in the NASA dataset are different in each project (Table 3.1).

Figure 6.2(c) shows the result for the RBM-generated weight sets in the AEEEM dataset. The figure shows that RBM generated different feature sets across projects in the AEEEM dataset. The explanation is similar to our observation during the correlation analysis. If we use one type of features, such as only the change features, Figure 6.2(d) shows that RBM generates the same feature sets across projects. However, if we use complexity code change features, we observe that RBM generates different feature sets across projects (Figure 6.2(e)). AE generated different features across projects (Figure 6.2(f), 6.2(g) and 6.2(h)).


From the discussion results, we can extract several possible explanations for the fact that RBM and AE significantly improve the variance of the performance across the unsupervised models, but not across the supervised models, and for why these feature reduction techniques perform well for the unsupervised models. As the studied neural network-based feature reduc-




(a) RBM-generated weight sets in the PROMISE dataset.




(b) RBM-generated weight sets in the NASA dataset.




(c) RBM-generated weight sets in the AEEEM dataset.




(d) RBM-generated weight sets using Change features in the AEEEM dataset.




(e) RBM-generated weight sets using complexity code change features in the AEEEM dataset.



(f) AE-generated weight sets in the PROMISE dataset.



(g) AE-generated weight sets in the NASA dataset.



(h) AE-generated weight sets using Change features in the AEEEM dataset.

**Figure 6.2** *k*-means clustering results with t-SNE for generated weight sets in the studied datasets. Each shape represents a project, and each color represents a cluster. Using color is necessary in this figure.

tion techniques appear to generate strongly correlated features for a project, these feature sets suffer from multicollinearity [7], which is known to negatively affect the performance of the supervised models. However, as the unsupervised defect prediction models do not need to be trained, these models are not affected by the multicollinearity problem. In addition, because RBM generates strongly correlated features for a project, the unsupervised models become much simpler, which seems to improve the variance of the performance across the unsupervised defect prediction models.

## 7. Threats to validity

### 7.1 External validity

With regards to the generalizability of our results, we applied our experiments to three publicly available datasets. These studied datasets (PROMISE, NASA and AEEEM) were all used in many prior defect prediction studies. The projects in these studied datasets span different domains, include both open source and industrial projects and contain different features. Future studies are necessary to investigate whether our results generalize to other projects.

In addition, we studied only a subset of the many existing feature reduction and selection techniques and defect prediction models. We carefully selected techniques and models that have been used before for defect prediction, and that have an implementation readily available. Without such an implementation, it is difficult and time-consuming to ensure that the implementation matches the one used in prior studies. Future studies are necessary to investigate whether our results apply to other feature reduction/selection techniques and defect prediction models.

### 7.2 Internal validity

In our experiments, we used AUC as a performance measure. AUC is a popular performance measure for defect prediction, as it does not require a threshold. However, different software project teams may have different objectives. Hence, future studies should investigate the impact of feature reduction techniques on other performance measures, while keeping in mind the possible pitfalls of studying threshold-dependent performance measures [73].

When using the out-of-sample bootstrap sampling, we encountered computational errors in two bootstrap samples. The errors occurred because two bootstrap samples violated requirements of the NB and SC models (e.g., a generated feature had a variance of zero which violates a requirement of the NB model). To mitigate this threat, we discarded these bootstrap samples and generated a new sample instead.

We provide all experimental scripts that we used in our study. <sup>(\*)</sup> This replication package allows researchers and practitioners to replicate our experiments and confirm our results.

---

(\*):<https://sailhome.cs.queensu.ca/replication/featred-vs-featsel-defectpred/>

## 8. Conclusion

In defect prediction, reducing the number of features is an important step when building defect prediction models [7,8,26,82]. Prior studies indicated that reducing the number of features avoids the problem of multicollinearity [7] and the curse of dimensionality [8]. Feature selection and reduction techniques help to reduce the number of features in a model. Feature selection techniques reduce the number of features in a model by selecting the most important ones, while feature reduction techniques reduce the number of features by creating new, combined features from the original features.

Prior work [9, 18] studied the impact of feature *selection* techniques on defect prediction models. Our work is the first large-scale study on the impact of feature *reduction* techniques on defect prediction models. In particular, we studied the impact of eight feature reduction techniques on five supervised and five unsupervised defect prediction models. In addition, we compared the impact of feature reduction techniques on defect prediction with the impact of the two best-performing feature selection techniques (according to prior work).

We studied the impact of feature reduction/selection techniques on defect prediction models along two dimensions:

1. The defect prediction performance (AUC) of the features that are generated by the feature reduction/selection techniques, to study whether feature reduction/selection techniques can improve the performance of defect prediction models.
2. The variance of the AUC across defect prediction models that use the features that are generated by feature reduction or selected by feature selection techniques. It is difficult to select the best performing model for each project, since the best model may change per project [11,77]. Hence, we studied whether feature reduction or selection techniques can relieve the burden for practitioners of having to choose the best performing defect prediction model for their data.

Below, we summarize the main recommendations that follow from our work.

**Recommendation 1: For the supervised defect prediction models, use the correlation-based (CFS) or consistency-based (ConFS) feature selection techniques.** Our experiments

in RQ3 show that, for the supervised models, CFS and ConFS outperform the feature reduction techniques (except feature agglomeration (FA)) and the original models. While FA has a similar performance, CFS and ConFS have a smaller performance variance. Hence, using CFS or ConFS in combination with a supervised defect prediction model allows practitioners to improve the performance of their defect prediction models, while making the choice for a particular defect prediction model easier as well.

**Recommendation 2: For the unsupervised defect prediction models, use a neural network-based technique (Restricted Boltzmann Machine (RBM) or autoencoder (AE)).** Our experiments in RQs 1 and 3 show that the RBM and AE feature reduction technique can significantly improve the performance of the unsupervised defect prediction models compared to the other feature reduction/selection techniques and the original models. In addition, we observed in RQs 2 and 3 that RBM and AE significantly improve the performance variance across the unsupervised models (except compared to the transfer component analyses (TCA and TCA+)). While the transfer component analyses (TCA and TCA+) have the smallest performance variance, they have a worse performance than RBM and AE. The effect size (Cliff's Delta) between the transfer component analyses and the neural network-based feature reduction techniques is negligible or small for the performance variance in favour of the transfer component analyses, but small (TCA) or large (TCA+) for the performance in favour of the neural network-based techniques. Hence, using a neural network-based feature reduction technique to preprocess the data of the unsupervised defect prediction models can improve both their performance and relieve the burden for practitioners of having to select the best-performing unsupervised defect prediction model for their project.

**Recommendation 3: If a project has diverse data, the neural network-based techniques (Restricted Boltzmann Machine (RBM) or autoencoder (AE)) are likely to improve its defect prediction performance and performance variance.** Our experiments showed that RBM and AE consistently improve the AUC and IQR of projects in the AEEEM dataset, for both supervised and unsupervised models. Practitioners should run PCA on their data to identify the diversity of their project's data (similar to what we did in Section 5.3.1). If the data turns out to be rich, RBM and AE are good options to improve the defect prediction performance and variance.

## Acknowledgment

This thesis would not have been possible without the support of many exceptional people to whom I am grateful.

First of all, I would like to thank my supervisor Dr. Osamu Mizuno for his supervision and continuous support throughout this journey. His technical and editorial advice is fundamental of this thesis. He also gave me amazing opportunities to discuss with the professional researchers all over the world. I learned so much from the discussion. He prepared the research environment that made the research faster and faster. I cannot thank you enough for all you have done.

I was a very lucky man to work and collaborate with some of the brightest researchers. I would like to thank all of my collaborators, Dr. Cor-Paul Bezemer, Dr. Yasutaka Kamei, Dr. Ahmed E. Hassan, and Dr. Eun-Hye Choi. They shared their experiences, offered a lot of advices, and taught writing techniques. The collaboration was invaluable to proceed and polish my work.

I have had the privilege of working with amazing groups at the Software Engineering Laboratory (SEL), the Software Analysis and Intelligence Lab (SAIL), and the Software Analytics Research Group. Each member has had an impact on me and my work. I am grateful to all members.

The work in this thesis was published on Empirical Software Engineering. I would like to thank all reviewers and editor. They gave me many insightful comments.

Without the support of my family, this thesis would not have been possible. I am particularly grateful to my parents.

## References

- [1] G. Tassej, The economic impacts of inadequate infrastructure for software testing, National Institute of Standards and Technology, 2002.
- [2] V.R. Basili, L.C. Briand, and W.L. Melo, “A validation of object-oriented design metrics as quality indicators,” *IEEE Transactions on Software Engineering (TSE)*, vol.22, no.10, pp.751–761, 1996.
- [3] S. Kim, T. Zimmermann, E.J. Whitehead Jr, and A. Zeller, “Predicting faults from cached history,” *Proceedings of the 29th International Conference on Software Engineering (ICSE)*, pp.489–498, IEEE Computer Society, 2007.
- [4] R. Moser, W. Pedrycz, and G. Succi, “A comparative analysis of the efficiency of change metrics and static code attributes for defect prediction,” *Proceedings of the 30th International Conference on Software Engineering (ICSE)*, pp.181–190, IEEE, 2008.
- [5] A.E. Hassan, “Predicting faults using the complexity of code changes,” *Proceedings of the 31st International Conference on Software Engineering (ICSE)*, pp.78–88, IEEE Computer Society, 2009.
- [6] M. D’Ambros, M. Lanza, and R. Robbes, “An extensive comparison of bug prediction approaches,” *Proceedings of the 7th International Conference on Mining Software Repositories (MSR)*, pp.31–41, IEEE, 2010.
- [7] D.E. Farrar and R.R. Glauber, “Multicollinearity in regression analysis: the problem revisited,” *The Review of Economic and Statistics*, vol.49, no.1, pp.92–107, 1967.
- [8] R. Bellman, *Dynamic programming*, Princeton University Press, 1957.
- [9] B. Ghotra, S. McIntosh, and A.E. Hassan, “A large-scale study of the impact of feature selection techniques on defect classification models,” *Proceedings of the 14th International Conference on Mining Software Repositories (MSR)*, pp.146–157, IEEE Press, 2017.
- [10] V.U.B. Challagulla, F.B. Bastani, I.-L. Yen, and R.A. Paul, “Empirical assessment of machine learning based software defect prediction techniques,” *International Journal on Artificial Intelligence Tools*, vol.17, no.02, pp.389–400, 2008.



- [11] A.R. Gray and S.G. Macdonell, “Software metrics data analysis—exploring the relative performance of some commonly used modeling techniques,” *Empirical Software Engineering*, vol.4, no.4, pp.297–316, 1999.
- [12] L. Guo, B. Cukic, and H. Singh, “Predicting fault prone modules by the dempster-shafer belief networks,” *Proceedings of the 18th International Conference on Automated Software Engineering (ASE)*, pp.249–252, IEEE, 2003.
- [13] S. Herbold, “Training data selection for cross-project defect prediction,” *Proceedings of the 9th International Conference on Predictive Models in Software Engineering*, p.6, ACM, 2013.
- [14] M. Jureczko and L. Madeyski, “Towards identifying software project clusters with regard to defect prediction,” *Proceedings of the 6th International Conference on Predictive Models in Software Engineering*, p.9, ACM, 2010.
- [15] Y. Kamei, T. Fukushima, S. McIntosh, K. Yamashita, N. Ubayashi, and A.E. Hassan, “Studying just-in-time defect prediction using cross-project models,” *Empirical Software Engineering*, vol.21, no.5, pp.2072–2106, 2016.
- [16] J. Nam and S. Kim, “CLAMI: Defect prediction on unlabeled datasets,” *Proceedings of the 30th International Conference on Automated Software Engineering (ASE)*, pp.452–463, IEEE, 2015.
- [17] J. Nam and S. Kim, “Heterogeneous defect prediction,” *Proceedings of the 10th Joint Meeting on Foundations of Software Engineering (FSE)*, pp.508–519, ACM, 2015.
- [18] Z. Xu, J. Liu, Z. Yang, G. An, and X. Jia, “The impact of feature selection on defect prediction performance: An empirical comparison,” *Proceedings of the 27th International Symposium on Software Reliability Engineering (ISSRE)*, pp.309–320, IEEE, 2016.
- [19] T. Zimmermann, N. Nagappan, H. Gall, E. Giger, and B. Murphy, “Cross-project defect prediction: a large scale experiment on data vs. domain vs. process,” *Proceedings of the 7th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on The Foundations of Software Engineering (ESEC-FSE)*, pp.91–100, ACM, 2009.

- [20] K. Gao, T.M. Khoshgoftaar, H. Wang, and N. Seliya, "Choosing software metrics for defect prediction: an investigation on feature selection techniques," *Software: Practice and Experience*, vol.41, no.5, pp.579–606, 2011.
- [21] S.S. Rathore and A. Gupta, "A comparative study of feature-ranking and feature-subset selection techniques for improved fault prediction," *Proceedings of the 7th India Software Engineering Conference*, p.7, ACM, 2014.
- [22] D. Rodriguez, R. Ruiz, J. Cuadrado-Gallego, J. Aguilar-Ruiz, and M. Garre, "Attribute selection in software engineering datasets for detecting fault modules," *Proceedings of the 2007 EUROMICRO Conference on Software Engineering and Advanced Applications*, pp.418–423, IEEE, 2007.
- [23] M.A. Hall and G. Holmes, "Benchmarking attribute selection techniques for discrete class data mining," *IEEE Transactions on Knowledge and Data Engineering*, vol.15, no.6, pp.1437–1447, 2003.
- [24] K. Muthukumaran, A. Rallapalli, and N. Murthy, "Impact of feature selection techniques on bug prediction models," *Proceedings of the 8th India Software Engineering Conference*, pp.120–129, ACM, 2015.
- [25] D. Rodríguez, R. Ruiz, J. Cuadrado-Gallego, and J. Aguilar-Ruiz, "Detecting fault modules applying feature selection to classifiers," *Proceedings of the 2007 International Conference on Information Reuse and Integration*, pp.667–672, IEEE, 2007.
- [26] S. Shivaji, E.J. Whitehead, R. Akella, and S. Kim, "Reducing features to improve code change-based bug prediction," *IEEE Transactions on Software Engineering (TSE)*, vol.39, no.4, pp.552–569, 2013.
- [27] J. Ren, K. Qin, Y. Ma, and G. Luo, "On software defect prediction using machine learning," *Journal of Applied Mathematics*, vol.2014, pp.1–13, 2014.
- [28] J. Nam, S.J. Pan, and S. Kim, "Transfer defect learning," *Proceedings of the 2013 International Conference on Software Engineering (ICSE)*, pp.382–391, IEEE Press, 2013.
- [29] C. Faloutsos and K.-I. Lin, "FastMap: A fast algorithm for indexing, data-mining and visualization of traditional and multimedia datasets," *Proceedings of the ACM SIGMOD*

- International Conference on Management of Data, pp.163–174, ACM, 1995.
- [30] L. Rokach and O. Maimon, “Clustering methods,” *Data mining and knowledge discovery handbook*, pp.321–352, Springer, 2005.
- [31] E. Bingham and H. Mannila, “Random projection in dimensionality reduction: applications to image and text data,” *Proceedings of the 7th International Conference on Knowledge Discovery and Data Mining*, pp.245–250, ACM, 2001.
- [32] S.J. Pan, I.W. Tsang, J.T. Kwok, and Q. Yang, “Domain adaptation via transfer component analysis,” *IEEE Transactions on Neural Networks*, vol.22, no.2, pp.199–210, 2011.
- [33] P. Smolensky, “Information processing in dynamical systems: Foundations of harmony theory,” *Technical report, DTIC Document*, 1986.
- [34] G.E. Hinton and R.R. Salakhutdinov, “Reducing the dimensionality of data with neural networks,” *Science*, vol.313, no.5786, pp.504–507, 2006.
- [35] M. Shepperd, Q. Song, Z. Sun, and C. Mair, “Data quality: Some comments on the NASA software defect datasets,” *IEEE Transactions on Software Engineering (TSE)*, vol.39, no.9, pp.1208–1215, 2013.
- [36] M. Kondo, C.-P. Bezemer, Y. Kamei, A.E. Hassan, and O. Mizuno, “The impact of feature reduction techniques on defect prediction models,” *Empirical Software Engineering*, p.To Appear, 2019.
- [37] I. Arora, V. Tatarwal, and A. Saha, “Open issues in software defect prediction,” *Procedia Computer Science*, vol.46, pp.906–912, 2015.
- [38] J. Nam, “Survey on software defect prediction,” *HKUST PhD Qualifying Examination, Department of Computer Science and Engineering, The Hong Kong University of Science and Technology, Tech. Rep*, pp.1–34, 2014.
- [39] Z. He, F. Shu, Y. Yang, M. Li, and Q. Wang, “An investigation on the feasibility of cross-project defect prediction,” *Automated Software Engineering*, vol.19, no.2, pp.167–199, 2012.
- [40] J. Nam, W. Fu, S. Kim, T. Menzies, and L. Tan, “Heterogeneous defect prediction,” *IEEE Transactions on Software Engineering*, pp.874–896, 2017.

- [41] F. Zhang, Q. Zheng, Y. Zou, and A.E. Hassan, "Cross-project defect prediction using a connectivity-based unsupervised classifier," Proceedings of the 38th International Conference on Software Engineering (ICSE), pp.309–320, ACM, 2016.
- [42] G. Abaei, Z. Rezaei, and A. Selamat, "Fault prediction by utilizing self-organizing map and threshold," Proceedings of the International Conference on Control System, Computing and Engineering (ICCSCE), pp.465–470, IEEE, 2013.
- [43] P.S. Bishnu and V. Bhattacharjee, "Software fault prediction using quad tree-based k-means clustering algorithm," IEEE Transactions on Knowledge and Data Engineering, vol.24, no.6, pp.1146–1150, 2012.
- [44] B. Yang, Q. Yin, S. Xu, and P. Guo, "Software quality prediction using affinity propagation algorithm," Proceedings of the International Joint Conference on Neural Networks, pp.1891–1896, IEEE, 2008.
- [45] S. Zhong, T.M. Khoshgoftaar, and N. Seliya, "Unsupervised learning for expert-based software quality estimation.," HASE, pp.149–155, Citeseer, 2004.
- [46] T. Menzies, J. Greenwald, and A. Frank, "Data mining static code attributes to learn defect predictors," IEEE Transactions on Software Engineering (TSE), vol.33, no.1, pp.2–13, 2007.
- [47] N. Nagappan, T. Ball, and A. Zeller, "Mining metrics to predict component failures," Proceedings of the 28th International Conference on Software Engineering (ICSE), pp.452–461, ACM, 2006.
- [48] D.E. Neumann, "An enhanced neural network technique for software risk analysis," IEEE Transactions on Software Engineering (TSE), vol.28, no.9, pp.904–912, 2002.
- [49] Z.M. Hira and D.F. Gillies, "A review of feature selection and feature extraction methods applied on microarray data," Advances in Bioinformatics, vol.2015, pp.1–13, 2015.
- [50] T. Menzies, D. Owen, and J. Richardson, "The strangest thing about software," Computer, vol.40, no.1, pp.54–60, 2007.
- [51] F. Peters, T. Menzies, L. Gong, and H. Zhang, "Balancing privacy and utility in cross-company defect prediction," IEEE Transactions on Software Engineering, vol.39, no.8,

pp.1054–1068, 2013.

- [52] S.R. Chidamber and C.F. Kemerer, “A metrics suite for object oriented design,” *IEEE Transactions on Software Engineering (TSE)*, vol.20, no.6, pp.476–493, 1994.
- [53] T.J. McCabe, “A complexity measure,” *IEEE Transactions on Software Engineering (TSE)*, vol.SE-2, no.4, pp.308–320, 1976.
- [54] M.H. Halstead, *Elements of software science*, vol.7, Elsevier New York, 1977.
- [55] J. Petrić, D. Bowes, T. Hall, B. Christianson, and N. Baddoo, “The jinx on the NASA software defect data sets,” *Proceedings of the 20th International Conference on Evaluation and Assessment in Software Engineering*, pp.1–5, ACM, 2016.
- [56] J.H. McDonald, *Handbook of Biological Statistics (3rd ed.)*, Sparky House Publishing, Baltimore, Maryland., 2014.
- [57] J. Han and C. Moraga, “The influence of the sigmoid function parameters on the speed of backpropagation learning,” *Proceedings of the International Workshop on Artificial Neural Networks*, pp.195–201, Springer, 1995.
- [58] R. Quinlan, *C4.5: Programs for Machine Learning*, Morgan Kaufmann Publishers, 1993.
- [59] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I.H. Witten, “The WEKA data mining software: an update,” *ACM SIGKDD explorations newsletter*, vol.11, no.1, pp.10–18, 2009.
- [60] T.K. Ho, “Random decision forests,” *Proceedings of the 3rd International Conference on Document Analysis and Recognition*, vol.1, pp.278–282, IEEE, 1995.
- [61] H. Zhang, “The optimality of Naive Bayes,” *Proceedings of 17th International Florida Artificial Intelligence Research Society Conference*, pp.562–567, AAAI Press, 2004.
- [62] N. Landwehr, M. Hall, and E. Frank, “Logistic model trees,” *Machine Learning*, vol.59, no.1, pp.161–205, 2005.
- [63] M. Kuhn, “Caret: classification and regression training,” 2015.
- [64] C. Tantithamthavorn, S. McIntosh, A.E. Hassan, and K. Matsumoto, “Automated parameter optimization of classification techniques for defect prediction models,” *Proceedings of*

- the 38th International Conference on Software Engineering (ICSE), pp.321–332, ACM, 2016.
- [65] U. von Luxburg, “A tutorial on spectral clustering,” *Statistics and computing*, vol.17, no.4, pp.395–416, 2007.
- [66] J.A. Hartigan and M.A. Wong, “Algorithm as 136: A k-means clustering algorithm,” *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, vol.28, no.1, pp.100–108, 1979.
- [67] L. Kaufman and P.J. Rousseeuw, *Finding groups in data: an introduction to cluster analysis*, vol.344, John Wiley & Sons, 2009.
- [68] J.C. Dunn, “A fuzzy relative of the isodata process and its use in detecting compact well-separated clusters,” *J. Cybernet*, vol.3, pp.32–57, 1973.
- [69] T. Martinetz and K. Schulten, “A “neural-gas” network learns topologies,” *Artificial Neural Networks*, vol.1, pp.397–402, 01 1991.
- [70] T. Kohonen, “The self-organizing map,” *Proceedings of the IEEE*, vol.78, no.9, pp.1464–1480, 1990.
- [71] M.A. Hall, “Correlation-based feature selection for machine learning,” PhD thesis, University of Waikato Hamilton, 1999.
- [72] M. Dash and H. Liu, “Consistency-based search in feature selection,” *Artificial intelligence*, vol.151, no.1, pp.155–176, 2003.
- [73] C. Tantithamthavorn and A.E. Hassan, “An experience report on defect modelling in practice: Pitfalls and challenges,” *Proceedings of the 40th International Conference on Software Engineering: Software Engineering in Practice Track (ICSE-SEIP)*, pp.286–295, ACM, 2018.
- [74] “Scikit learn: Minmaxscaler,” <http://scikit-learn.org/stable/modules/generated/sklearn.preprocessing.MinMaxScaler.html>. [Online; accessed 1-July-2018].
- [75] C. Tantithamthavorn, S. McIntosh, A.E. Hassan, and K. Matsumoto, “An empirical comparison of model validation techniques for defect prediction models,” *IEEE Transactions*

- on Software Engineering (TSE), vol.43, no.1, pp.1–18, 2017.
- [76] J. Cohen, “Statistical power analysis for the behavioral sciences,” 1988.
- [77] B. Ghotra, S. McIntosh, and A.E. Hassan, “Revisiting the impact of classification techniques on the performance of defect prediction models,” Proceedings of the 37th International Conference on Software Engineering (ICSE), pp.789–800, IEEE Press, 2015.
- [78] E. Shihab, Z.M. Jiang, W.M. Ibrahim, B. Adams, and A.E. Hassan, “Understanding the impact of code and process metrics on post-release defects: A case study on the Eclipse project,” Proceedings of the International Symposium on Empirical Software Engineering and Measurement (ESEM), pp.4:1–4:10, ACM, 2010.
- [79] D. Zwillinger and S. Kokoska, CRC standard probability and statistics tables and formulae, Crc Press, 1999.
- [80] L. van der Maaten and G. Hinton, “Visualizing data using t-SNE,” Journal of Machine Learning Research, vol.9, no.Nov, pp.2579–2605, 2008.
- [81] L. van der Maaten, “Accelerating t-SNE using tree-based algorithms.,” Journal of Machine Learning Research, vol.15, no.1, pp.3221–3245, 2014.
- [82] E. Shihab, “Practical software quality prediction,” Proceedings of the 2014 International Conference on Software Maintenance and Evolution (ICSME), pp.639–644, IEEE, 2014.