# Comments on "Researcher Bias: The Use of Machine Learning in Software Defect Prediction"

Chakkrit Tantithamthavorn, *Student Member, IEEE*,
Shane McIntosh, *Member, IEEE*,
Ahmed E. Hassan, *Member, IEEE*,
and Kenichi Matsumoto, *Senior Member, IEEE*

**Abstract**—Shepperd et al. find that the reported performance of a defect prediction model shares a strong relationship with the group of researchers who construct the models. In this paper, we perform an alternative investigation of Shepperd et al.'s data. We observe that (a) research group shares a strong association with other explanatory variables (i.e., the dataset and metric families that are used to build a model); (b) the strong association among these explanatory variables makes it difficult to discern the impact of the research group on model performance; and (c) after mitigating the impact of this strong association, we find that the research group has a smaller impact than the metric family. These observations lead us to conclude that the relationship between the research group and the performance of a defect prediction model are more likely due to the tendency of researchers to reuse experimental components (e.g., datasets and metrics). We recommend that researchers experiment with a broader selection of datasets and metrics to combat any potential bias in their results.

**Index Terms**— Software quality assurance, defect prediction, researcher bias

◆

## 1 INTRODUCTION

RECENTLY, Shepperd et al. [1] study the extent to which the research group that performs a defect prediction study associates with the reported performance of defect prediction models. Through a meta-analysis of 42 primary studies, they find that the reported performance of a defect prediction model shares a strong relationship with the group of researchers who construct the models. Shepperd et al.'s findings raise several concerns about the current state of the defect prediction field. Indeed, their findings suggest that many published defect prediction studies are biased, and calls their validity into question.

In this paper, we perform an alternative investigation of Shepperd et al.'s data. More specifically, we set out to investigate (1) the strength of the association among the explanatory variables, e.g., research group and metric family (Section 2); (2) the interference that these associations introduce when interpreting the impact that explanatory variables have on the reported performance (Section 3); and (3) the impact that the explanatory variables have on the reported performance after we mitigate the strong associations among the explanatory variables (Section 4).

## 2 THE PRESENCE OF COLLINEARITY

We suspect that research groups are likely to reuse experimental components (e.g., datasets, metrics, and classifiers) in several studies. This tendency to reuse experimental components would introduce a strong association among the explanatory variables of

- C. Tantithamthavorn and K. Matsumoto are with the Graduate School of Information Science, Nara Institute of Science and Technology, Nara, Japan.
  E-mail: {chakkrit-t, matsumoto}@is.naist.jp.
- S. McIntosh is with the Department of Electrical and Computer Engineering, McGill University, Montreal, QC, Canada. E-mail: shane.mcintosh@mcgill.ca.
- A.E. Hassan is with the School of Computing, Queen's University, Kingston, ON, Canada. E-mail: ahmed@cs.queensu.ca.

Shepperd et al. [1]. To investigate our suspicion, we measure the strength of the association between each pair of the explanatory variables that are used by Shepperd et al. [1], i.e., ResearcherGroup, DatasetFamily, MetricFamily, and ClassifierFamily.

*Approach.* Since the explanatory variables are categorical, we first use a Pearson $\chi^2$ test [2] to check whether a statistically significant association exists between each pair of explanatory variables ($\alpha = 0.05$). Then, we compute Cramer's $V$ [3] to quantify the strength of the association between each pair of two categorical variables. The value of Cramer's $V$ ranges between 0 (no association) and 1 (strongest association). We use the convention of Rea et al. [4] for describing the magnitude of an association. To compute the Pearson's $\chi^2$ and Cramer's $V$ values, we use the implementation provided by the `assocstats` function of the `vcd` R package [5].

*Results. Research group shares a strong association with the dataset and metrics that are used.* Table 1 shows the Cramer's $V$ values and the p-value of the Pearson $\chi^2$ test for each pair of explanatory variables. The Cramer's $V$ values indicate that research group shares a strong association with the dataset and metrics that are used. Indeed, we find that 13 of the 23 research groups (57%) only experiment with one dataset family, where nine of them only use one NASA dataset, which contains only one family of software metrics (i.e., static metrics). Moreover, 39% of research groups only use the static metric family of the NASA dataset in several studies. The strong association among research groups, dataset, and metrics confirms our suspicion that researchers often reuse experimental components.

## 3 THE INTERFERENCE OF COLLINEARITY

The strong association among explanatory variables that we observe in Section 2 may introduce interference when one studies the impact that these explanatory variables have on the outcome [6], [7]. Furthermore, this interference among variables may cause impact analyses, such as ANOVA, to report spurious relationships that are dependent on the ordering of variables in the model formula. Indeed, ANOVA is a hierarchical model that first attributes as much variance as it can to the first variable before attributing residual variance to the second variable in the model formula [8]. If two variables share a strong association, the variable that appear first in the model formula will have the brunt of the variance associated with it. Hence, we set out to investigate the interference that is introduced by the strong association among explanatory variables.

*Approach.* To investigate this interference, we use a bootstrap analysis approach, which leverages aspects of statistical inference [9]. We first draw a bootstrap sample of size $N$ that is randomly drawn with replacement from an original dataset that is also of size $N$. We train linear regression models with the data of the bootstrap sample using the implementation provided by the `lm` function of the `stats` R package [10]. For each bootstrap sample, we train multi-way ANOVA models with all of the 24 possible ordering of the explanatory variables (e.g., ANOVA(ResearcherGroup $\times$ DatasetFamily $\times$ MetricFamily $\times$ ClassifierFamily) versus ANOVA(DatasetFamily $\times$ ResearcherGroup $\times$ MetricFamily $\times$ ClassifierFamily)). Following the prior study [1], we compute the partial $\eta^2$ values [11], which describe the proportion of the total variance that is attributed to an explanatory variable for each of the 24 models. We use the implementation provided by the `etasq` function of the `heplots` R package [12]. We repeat the experiment 1,000 times for each of the 24 models to produce a distribution of the partial $\eta^2$ values for each explanatory variable.

*Results. The strong association among the explanatory variables introduces interference when interpreting the impact that research group has on model performance.* Fig. 1 shows the distributions of the partial $\eta^2$ values for each explanatory variable when it appears at each position in the model formula. Each boxplot is derived from the models

TABLE 1
The Association Among Explanatory Variables

| Pair | Cramer's $V$ | Magnitude |
|---|---|---|
| ResearcherGroup & MetricFamily | 0.65*** | Strong |
| ResearcherGroup & DatasetFamily | 0.56*** | Relatively strong |
| MetricFamily & DatasetFamily | 0.55*** | Relatively strong |
| ResearcherGroup & ClassifierFamily | 0.54*** | Relatively strong |
| DatasetFamily & ClassifierFamily | 0.34*** | Moderate |
| MetricFamily & ClassifierFamily | 0.21*** | Moderate |

*Statistical significance of the Pearson $\chi^2$ test:$\circ p \geq .05$; * $p < .05$; ** $p < .01$; *** $p < .001$.*

of the 24 possible variable ordering combinations. The results show that there is a decreasing trend in partial eta-squared values when collinearity is not mitigated. Indeed, research group and dataset family tend to have the largest impact when they appear in earlier positions in the model formula, indicating that the impact that explanatory variables have on the outcome depends on the ordering of variables in the model formula. Moreover, we observe that research group and dataset family tend to have comparable partial $\eta^2$ values, indicating that the strong association introduces interference. In particular, a model with DatasetFamily in the first position and ResearcherGroup in the second position in the model formula would lead us to believe that DatasetFamily plays a bigger role than the ResearcherGroup. However, once we flip the positions of DatasetFamily and ResearcherGroup, we would reach a different conclusion, i.e., that ResearcherGroup plays a bigger role in the model. On the other hand, the moderate association between the metric and classifier families introduces much less interference.

## 4  MITIGATING COLLINEARITY

In the study of Shepperd et al., Table 13 is derived from a one-way ANOVA analysis ($y = x_n$) for each of the explanatory variables and Table 14 is derived from a multi-way ANOVA analysis ($y = x_1 \times \cdots \times x_n$). One of the main assumption of ANOVA analysis is that the explanatory variables must be independent. The prior sections show that such strong association among the explanatory variables (Section 2) introduces interference when we interpret the impact that the explanatory variables have on the outcome (Section 3). However, Shepperd et al.'s multi-way ANOVA analysis did not mitigate for the collinearity between the explanatory variables.

In this section, we set out to investigate the impact that the explanatory variables have on the reported performance after mitigating the interference that is introduced by strongly associated explanatory variables. Thus, our earlier analysis indicates that we cannot include all three of the explanatory variables in the same model. Instead, we should only include two of these three variables in our models. Hence, we opt to focus on one studied dataset in order to control the dataset family metric by holding it constant.

*Approach.* To mitigate the interference, we first select the Eclipse dataset family, which is the second-most popular dataset family in the studied dataset. We choose the Eclipse dataset family instead of selecting the NASA dataset family because Section 2 reveals that 39 percent of the research groups only use the static metric family of the NASA dataset in several studies. Overall, the Eclipse dataset family contains six metrics families, which are used by seven research groups who fit defect prediction model using six classifier families. While controlling for the dataset family metric, only research group, metric family, and classifier family will be included in the model formulas. Since the dataset family metric is now a constant, it is excluded from our model formula.

Since Table 1 shows that research group shares a strong association with the metric family, we build two different linear regression models by removing one of the two strongly associated variables, i.e., one model uses the research group and classifier
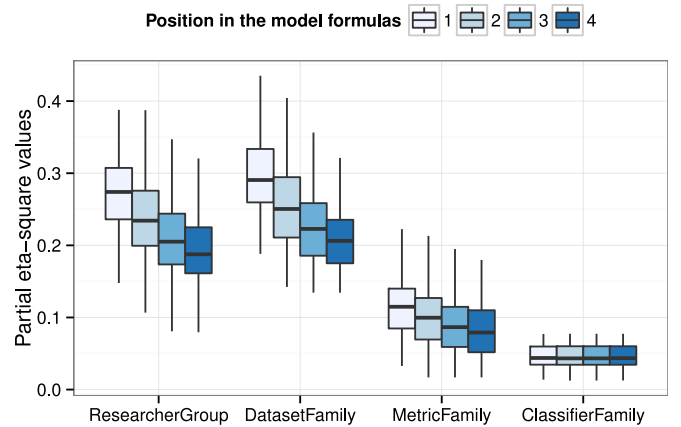


Fig. 1. Distribution of the partial $\eta^2$ values for each explanatory variable when it appears at each position in the model formula.

family variables, while another model uses the metric family and classifier family variables.

To confirm the absence of collinearity in the models, we perform a redundancy analysis [13] in order to detect redundant variables prior to constructing the models. We use the implementation provided by the `redun` function in the `rms` R package [14].

To assess the fit of our models, we compute the adjusted $R^2$ and the Akaike Information Criterion (AIC) [15]. The adjusted $R^2$ measures the amount of variance, while AIC measures the goodness-of-fit based on information entropy. In general, higher adjusted $R^2$ and lower AIC values correspond to a better fit of the model to the underlying data.

Finally, we perform a multi-way ANOVA analysis and compute the partial $\eta^2$ values. As suggested by Richardson et al. [11] and Mittas et al. [16], we use the convention of Cohen [17] for describing the effect size of the partial $\eta^2$ values — values below 0.01, 0.06, and 0.14 describe small, medium, and large effect sizes, respectively.

*Results. When we mitigate the interference of strongly associated explanatory variables, we find that the research group has a smaller impact than the metric family with respect to the Eclipse dataset family.* Table 2 shows partial $\eta^2$ values of the ANOVA analysis with respect to the Eclipse dataset family. The results show that the MetricFamily model, which achieves a higher adjusted $R^2$ and a lower AIC, tends to represent the underlying data better than the ResearcherGroup model. The redundancy analysis also confirms that there are no redundant variables in the MetricFamily model. Unlike Shepperd et al.'s earlier observations, our ANOVA analysis of the MetricFamily model shows that the choice of metrics that are used

TABLE 2
Partial $\eta^2$ Values of the Multi-Way ANOVA Analysis
with Respect to the Eclipse Dataset Family

| | ResearcherGroup Model | MetricFamily Model |
|---|---|---|
| Adjusted $R^2$ | 0.19 | 0.36 |
| AIC | −77 | −105 |
| | Partial $\eta^2$ | Partial $\eta^2$ |
| Research Group | 0.127 (medium) | † |
| Metric Family | † | 0.235 (large) |
| Classifier Family | 0.122 (medium) | 0.113 (medium) |
| Research Group: Classifier Family | 0.022 (small) | - |
| Metric Family: Classifier Family | - | 0.162 (large) |

†*Strong association variables.*

to build defect prediction models tends to have a large impact on the reported performance with respect to the Eclipse dataset family. Moreover, since the interference has been mitigated, the ANOVA results still hold when the explanatory variables are reordered.

## 5  CONCLUSIONS

The prior work of Shepperd et al. [1] suggests that the reported performance of a defect prediction model shares a strong relationship with the research group who conducted the study. This observation raises several concerns about the state of the defect prediction field. In this paper, we investigate (1) the strength of the association among the explanatory variables of Shepperd et al.'s study [1]; (2) the interference that these associations introduce when interpreting the impact of the explanatory variables on the reported performance; and (3) the impact that the explanatory variables have on the reported performance after we mitigate the interference introduced by strongly associated explanatory variables. We make the following observations:

- Research group shares a strong association with the dataset and metrics families that are used in building models, suggesting that researchers should experiment with a broader selection of datasets and metrics in order to maximize external validity.

- The strong association among explanatory variables introduces interference when interpreting the impact that research group has on the reported model performance, suggesting that researchers should carefully mitigate collinearity issues prior to analysis in order to maximize internal and construct validity.

- After mitigating the interference, we find that the research group has a smaller impact than metric family with respect to the Eclipse dataset family, suggesting that researchers should carefully examine the choice of metrics when building defect prediction models.

These observations lead us to conclude that the relationship between research groups and the performance of a defect prediction model have more to do with the tendency of researchers to reuse experimental components (e.g., datasets and metrics). Hence, a threat of bias exists if authors fixate on studying the same datasets with the same metrics. We recommend that research groups experiment with different datasets and metrics rather than relying entirely on reusing experimental components.

When adhering to our recommendation, researchers should be mindful of the inherent trade-off between maximizing internal and external validity in empirical research [18]. For example, maximizing external validity by studying a large corpus of datasets may raise threats to the internal validity of the study (i.e., the insights may be difficult to discern due to a broad selection of the studied systems). On the other hand, maximizing internal validity by focusing on a highly controlled experiment may raise threats to the external validity of the study (i.e., the insights may be too specific to the studied systems to generalize to other systems).

## REFERENCES

[1] M. Shepperd, D. Bowes, and T. Hall, "Researcher bias: The use of machine learning in software defect prediction," *IEEE Trans. Softw. Eng.*, vol. 40, no. 6, pp. 603–616, Jun. 2014.

[2] A. Agresti, *An introduction to categorical data analysis*. New York, NY, USA: Wiley, vol. 135, 1996.

[3] H. Cramér, *Mathematical Methods of Statistics*. Princeton, NJ, USA: Princeton Univ. Press, vol. 9, 1999.

[4] L. M. Rea and R. A. Parker, *Designing and Conducting Survey Research: A Comprehensive Guide*. Hoboken, NJ, USA: Wiley, 2014.

[5] D. Meyer, A. Zeileis, K. Hornik, F. Gerber, and M. Friendly, "VCD: Visualizing categorical data," 2015. [Online]. Available: http://CRAN.R-project.org/package=vcd.

[6] R. Grewal, J. a. Cote, and H. Baumgartner, "Multicollinearity and measurement error in structural equation models: Implications for Theory Testing," *Marketing Sci.*, vol. 23, no. 4, pp. 519–529, 2004.

[7] Y.-K. Tu, M. Kellett, V. Clerehugh, and M. S. Gilthorpe, "Problems of correlations between explanatory variables in multiple regression analyses in the dental literature." *Brit. Dental J.*, vol. 199, no. 7, pp. 457–461, 2005.

[8] B. R. Clifford, "Tests of hypotheses for unbalanced factorial designs under various regression/coding method combinations," *Educational Psychological Meas.*, no. 38, pp. 621–631, 1978.

[9] B. Efron and R. J. Tibshirani, *An Introduction to the Bootstrap*. Boston, MA, USA: Springer, 1993.

[10] R Core Team, "R: A language and environment for statistical computing," 2013. [Online]. Available: http://www.R-project.org/

[11] J. T. E. Richardson, "Eta squared and partial eta squared as measures of effect size in educational research," *Educational Res. Rev.*, vol. 6, no. 2, pp. 135–147, 2011.

[12] M. Friendly, "Heplots: Visualizing hypothesis tests in multivariate linear models," 2015. [Online]. Available: http://CRAN.R-project.org/package=heplots

[13] F. E. Harrell Jr., *Regression Modeling Strategies*, 1st ed. New York, NY, USA: Springer, 2002.

[14] F. E. Harrell Jr., "RMS: Regression modeling strategies," 2015. [Online]. Available: http://CRAN.R-project.org/package=rms

[15] H. Akaike, "A new look at the statistical model identification," *IEEE Trans. Autom. Control*, vol. 19, no. 6, pp. 716–723, Dec. 1974.

[16] N. Mittas and L. Angelis, "Ranking and clustering software cost estimation models through a multiple comparisons algorithm," *IEEE Trans. Softw. Eng.*, vol. 39, no. 4, pp. 537–551, Apr. 2013.

[17] J. Cohen, *Statistical Power Analysis for the Behavioral Sciences*. New York, NY, USA: Academic, 1988.

[18] J. Siegmund, N. Siegmund, and S. Apel, "Views on internal and external validity in empirical software engineering," in *Proc. Int. Conf. Softw. Eng.*, 2015, pp. 1276–1304.

---

1. http://sailhome.cs.queensu.ca/replication/researcher_bias_comments/
2. https://www.computecanada.ca/
3. http://cac.queensu.ca/