Metric Learning Improves the Ability of Combinatorial Coverage Metrics to Anticipate Classification Error

Tyler Cody
National Security Institute
Virginia Tech
Arlington, VA, USA

Laura Freeman
National Security Institute
Virginia Tech
Arlington, VA, USA

Abstract—Machine learning models are increasingly used in practice. However, many machine learning methods are sensitive to test or operational data that is dissimilar to training data. Out-of-distribution (OOD) data is known to increase the probability of error and research into metrics that identify what dissimilarities in data affect model performance is on-going. Recently, combinatorial coverage metrics have been explored in the literature as an alternative to distribution-based metrics. Results show that coverage metrics can correlate with classification error. However, other results show that the utility of coverage metrics is highly dataset-dependent. In this paper, we show that this dataset-dependence can be alleviated with metric learning, a machine learning technique for learning latent spaces where data from different classes is further apart. In a study of 6 open-source datasets, we find that metric learning increased the difference between set-difference coverage metrics (SDCCMs) calculated on correctly and incorrectly classified data, thereby demonstrating that metric learning improves the ability of SDCCMs to anticipate classification error. Paired t-tests validate the statistical significance of our findings. Overall, we conclude that metric learning improves the ability of coverage metrics to anticipate classifier error and identify when OOD data is likely to degrade model performance.

Index Terms—combinatorial coverage, combinatorial testing, machine learning, out-of-distribution (OOD)

I. INTRODUCTION

Machine learning is an increasingly dominant part of software systems and, thus, is of increasing interest in the software testing community. Compared to the standard statistical performance analysis on independent, identically-distributed data that machine learning practitioners and researchers typically use to test machine learning [1], the perspective of the software testing community has a broader scope. The learned models produced by machine learning algorithms are function approximations, and, as such, can be treated with the general test methods developed for system functions [2]. A key set of general test methods identified for software functions is combinatorial interaction testing [3].

Combinatorial interaction testing is typically concerned with the interaction of various system components and typically applied during system integration to determine the subset of interactions among components that cause the majority of failures. It has been shown reported that most software faults can be attributed to 2-, 3-, or 4-way interactions among software components [4]. Recently, combinatorial interaction testing and related concepts have been extended to machine learning. A key concept in related machine learning literature is combinatorial coverage [3].

Combinatorial coverage measures the number of interactions present in a dataset against the total number of possible interactions. Early applications of combinatorial coverage to machine learning include the use of coverage within deep learning algorithms [5]–[7]. These early works had mixed results, and faced heavy criticism [8]. Central to critiques was the idea that combinatorial coverage was concerned with the total number of possible interactions, or, otherwise put, the entire universe of interactions. Most of these interactions are not natural, and many can be even be considered adversarial [8].

While coverage metrics faced difficulties within machine learning algorithms, their application to the testing of machine learning algorithms grew in the literature. Early examples include use cases in explainability [9], autonomous vehicles [10], and test set construction [11]. Building on this early work, Lanus et. al proposed set-difference combinatorial coverage metrics (SDCCMs) which, in contrast to traditional coverage metrics [3], did not calculate coverage with respect to the entire universe of interactions, but rather with respect to another dataset [12]. SDCCMs can be used to compare the interactions seen in training, to those seen in testing, to those seen in operation—all relative to each other, not relative to all possible interactions. Cody et. al use SDCCMs to directly refute previous critiques of the use of combinatorial coverage in machine learning using MNIST image data [13].

While previous work has affirmed the utility of combinatorial coverage concepts to machine learning, their use has also been shown to be highly dataset-dependent [14]. In this work, we propose that metric learning be used to address dataset-dependence. Metric learning is a subfield of machine learning concerned with learning latent spaces where data from different classes is further apart than in their original space. In a study of 6 open-source datasets from a variety of domains and with a variety of sizes, we provide evidence that

metric learning improves the ability of SDCCMs to anticipate classification error. In doing so, we also provide supporting evidence for existing claims in the literature that SDCCMs are correlated to classification error [12], [13].

This paper is structured as follows. First, related work is reviewed. Combinatorial coverage and set-difference combinatorial coverage (SDCC) are defined as given by [12]. Then, background is provided on metric learning and details are provided for the metric learning methods used in this paper. The method of analysis is described, results are presented with accompanying discussion, and, lastly, conclusions are drawn.

II. PRELIMINARIES

Notions of combinatorial coverage have been applied to deep learning [5]–[7]. These applications have been critiqued [8]. Broader use of coverage concepts in machine learning include explainability [9], [15], autonomous vehicles [10], [16], test set construction [11], fairness testing [17], physical unclonable functions [9], multi-domain operations [18], and active learning [19]. Many works focus on traditional measures of coverage [3], however recent work proposes new measures for applications of coverage to machine learning based on set-difference [12].

A. Coverage Formalism

In combinatorial interaction testing for machine learning, *factors* are the features $\mathcal{X}$ and classes $\mathcal{Y}$ and the *values* of those factors are their events. Note, in combinatorial interaction testing, continuous-valued factors must be discretized to a finite set of values. A t-way *value combination* is a t-tuple of (factor, value) pairs. If there are k factors in $\mathcal{X} \times \mathcal{Y}$, then each element $(x, y) \in \mathcal{X} \times \mathcal{Y}$ contains $\binom{k}{t}$ t-way value combinations. Value combinations are a formal description of interactions.

Combinatorial coverage, sometimes termed total t-way coverage, measures the proportion of valid t-way value combinations that appear in a set [3]. Value combinations that appear are considered *covered*. Combinatorial coverage is defined in the following.

Definition 1. t-way Combinatorial Coverage.
Consider a universe with k factors such that $\mathcal{U}$ is the set of all valid k-way value combinations. Let $\mathcal{U}^t$ be the set of valid t-way combinations. Given a set of data $D \subseteq \mathcal{U}$, let D^t define the set of t-way value combinations appearing in D. The t-way combinatorial coverage of D is

$$CC^t(D) = \frac{|D^t|}{|\mathcal{U}^t|},$$

where $|D|$ denotes the cardinality of D.

SDCC is alternatively measures the proportion of valid t-way value combinations that appear in one set relative to another set. SDCC is defined in the following [12].

Definition 2. t-way Set Difference Combinatorial Coverage.
Let D_A and D_B be sets of data, and $D_A{}^t$ and $D_B{}^t$ be the corresponding t-way sets of data. The set difference $D_B{}^t \backslash D_A{}^t$ gives the value combinations that are in $D_B{}^t$ but that are not in $D_S{}^t$. The t-way set difference combinatorial coverage is

$$SDCC^t(D_B, D_A) = \frac{|D_B{}^t \backslash D_A{}^t|}{|D_B{}^t|}.$$

$SDCC^t$ is bounded $[0, 1]$ where 1 indicates no overlap, i.e., $D_B{}^t \cap D_A{}^t = \emptyset$, and 0 indicates that $D_B{}^t \subseteq D_A{}^t$. In other words, if D_B is the testing data and D_A is the training data, 0 indicates that all testing combinations are present in the training data and 1 indicates that none are present. Note, we use the term SDCCMs to refer to SDCC over a set of t-values.

III. METRIC LEARNING BACKGROUND

Metric learning is generally concerned with learning distance functions [20]. This typically involves learning a transformation to a latent space where similar data are closer together and dissimilar data are farther apart. As such, the latent spaces learned by metric learning induce sparsity in the input space that is both useful for making predictions (by making data more linearly separable). This may also be useful for coverage as the more separated data is easier to discretize [21]. Although deep metric learning methods are an active area of interest [22], [23], critiques suggest that, so far, deep metric learning offers little improvement over classical methods in cases where both deep and classical methods can be applied [24].

We use three of the most prevalent classical methods: Neighborhood Component Analysis (NCA) [25], Metric Learning for Kernel Regression (MLKR) [26], and Large Margin Nearest Neighbor (LMNN) [27].

- NCA is a variation on k-nearest neighbors (KNN) classification that directly maximizes a variant of leave-one-out performance.
- MLKR learns a distance function by directly minimizing the leave-one-out regression error.
- LMNN is a variation on KNN classification that learns a Mahalanobis distance.

These three methods are supervised metric learning methods because they use class labels during training.

IV. DATA

In this paper, 6 open-source datasets are used. All datasets can be retrieved from the University of California Irvine (UCI) Machine Learning Repository [28]. The datasets are described in Table I and are: the "Wine Data Set" (Wine), the "Rice (Cammeo and Osmancik) Data Set" (Rice), the "Yeast Data Set" (Yeast), the "Car Evaluation Data Set" (Car), the "Breast Cancer Wisconsin (Diagnostic) Data Set" (Cancer), and the "Balance Scale Data Set" (Balance). The datasets vary in size, number of classes, number of features, and feature type. The Cancer dataset is reduced from its original feature space to only the 'mean' features. All other datasets are unaltered.

TABLE I
Dataset Information

Dataset	Sample Size	Classes	Features	Feature Type
Wine	178	3	13	Continuous
Rice	3810	2	7	Continuous
Yeast	1484	10	8	Continuous
Car	1728	4	6	Discrete
Cancer	569	2	10	Continuous
Balance	625	3	4	Discrete

TABLE II
Classification Accuracy

n = 10 Dataset	DT (Mean, Std)	SVM (Mean, std)	KNN (Mean, Std)	Random 1/No. Classes
Wine	(0.85, 0.07)	(0.61, 0.06)	(0.72, 0.08)	0.33
Rice	(0.92, 0.01)	(0.88, 0.14)	(0.88, 0.02)	0.50
Yeast	(0.48, 0.04)	(0.59, 0.03)	(0.56, 0.03)	0.10
Car	(0.78, 0.02)	(0.93, 0.02)	(0.90, 0.02)	0.25
Cancer	(0.92, 0.02)	(0.89, 0.03)	(0.89, 0.03)	0.50
Balance	(0.64, 0.04)	(0.90, 0.03)	(0.82, 0.03)	0.33

V. Method

In this paper, we investigate whether metric learning improves the ability of SDCCMs to anticipate classification error. To assess the ability to anticipate error, we measure the difference between SDCCMs calculated (1) between training data and correctly classified data and (2) between training data and incorrectly classified data. We suppose that the larger the difference, the better SDCCMs perform as a metrics for anticipating classification error. To conduct our analysis, the data must be prepared.

The data is prepared for coverage analysis as follows. The NCA, MLKR, and LMNN algorithms are used to learn transformations to latent spaces using all available data. The metric learning methods use the default settings of metric-learn [29]. Each dimension of the latent spaces is used as a feature. The number of dimensions is equal to the number of features in the original data. Then, the NCA, MLKR, LMNN, and original spaces are discretized. Note, the Car and Balance datasets have discrete features and thus the discretization is not applied to their original spaces. The discretization process treats each feature separately. Each feature's data is clustered using the standard k-means clustering algorithm from scikit-learn with k=5 bins [30].

After transformation and discretization, the data is ready for coverage analysis. The coverage analysis is conducted as follows. For each dataset, 10 folds are created with 80%-20% train-test splits using random sampling. For each fold, a decision tree (DT), support vector machine (SVM), and KNN classifier is trained and tested. The classifiers use the default scikit-learn settings [30], except the maximum depth is set to 2 for the DT classifier. The correctly and incorrectly classified test data are separated. Then, the SDCCMs are calculated over the discretized NCA, MLKR, LMNN, and original features (i.e., not the labels) (1) between the training data and the correctly classified data and (2) between the training data and the incorrectly classified data. In both cases, the SDCC is used to measure the ratio of the number of interactions that appear in the training data that do not appear in test data to the total number of interactions that appear in the training data. In previous literature [13] and during initial testing, this direction of set-difference showed stronger correlation with error than the ratio of interactions appearing in the test data that do not appear in the training data to the total number of interactions in the test data. The difference between SDCCMs calculated in (1) and (2) is assessed for $t = (2, 3, 4)$.

VI. Results

The results of the coverage analysis are examined visually and statistically to assess whether metric learning increases the ability of SDCCMs to anticipate classification error. The accuracy scores of the classifiers are summarized in Table II. Figures 1 and 2 visualize the results. The rows of the figures correspond to the datasets. The columns of the figures correspond to the DT, SVM, and KNN classifiers, respectively, from left-to-right.

Within each subplot, the y-axis corresponds to the SDCC for the incorrectly classified data subtracted from the SDCC for the correctly classified test data. Higher values along the y-axis indicate a higher ability to anticipate classification error. The x-axis corresponds to the t-values used in the SDCC calculations, and boxplots are grouped according to the t-values. The different colors of the boxplots correspond to the discretized original, NCA, MLKR, and LMNN spaces, respectively, from left-to-right. The center line of the boxes indicates the median, the upper and lower borders of the boxes indicate the upper and lower quartiles, and the whiskers indicate the rest of the distribution. Outliers are depicted as black diamonds. Note, these outliers are identified by their distance from the inter-quartile range, and are not removed in later statistical analysis.

The general trends in the figures can be described as follows. In Figure 1, the Wine dataset shows outperformance by metric learning for nearly all classifiers, t-values, and metric learning methods. The Yeast dataset, in contrast, depicts mixed results, with at least one metric learning method outperforming in many cases, but no cases of all metric learning methods outperforming. Lastly, the Rice dataset shows the opposite results of the Wine dataset, with the original space outperforming for nearly all cases. This dataset-to-dataset variance in results is expected, and has been noted in previous literature [14]. The trends in Figure 2 are more nuanced. For the Car dataset, metric learning spaces tend to outperform the original space more for lower values of t, while for the Cancer dataset, metric learning outperforms for higher values of t. The Balance dataset shows both of these trends, with, in general, more outperformance occurring for $t = 2$ and $t = 4$ than for $t = 3$.

In summary, the figures show varied results, however, for many cases, metric learning clearly increases the ability of SDCCMs to anticipate classification error. From visual inspection of Figures 1 and 2, the largest dependency in the success of metric learning appears to be on the dataset, followed by
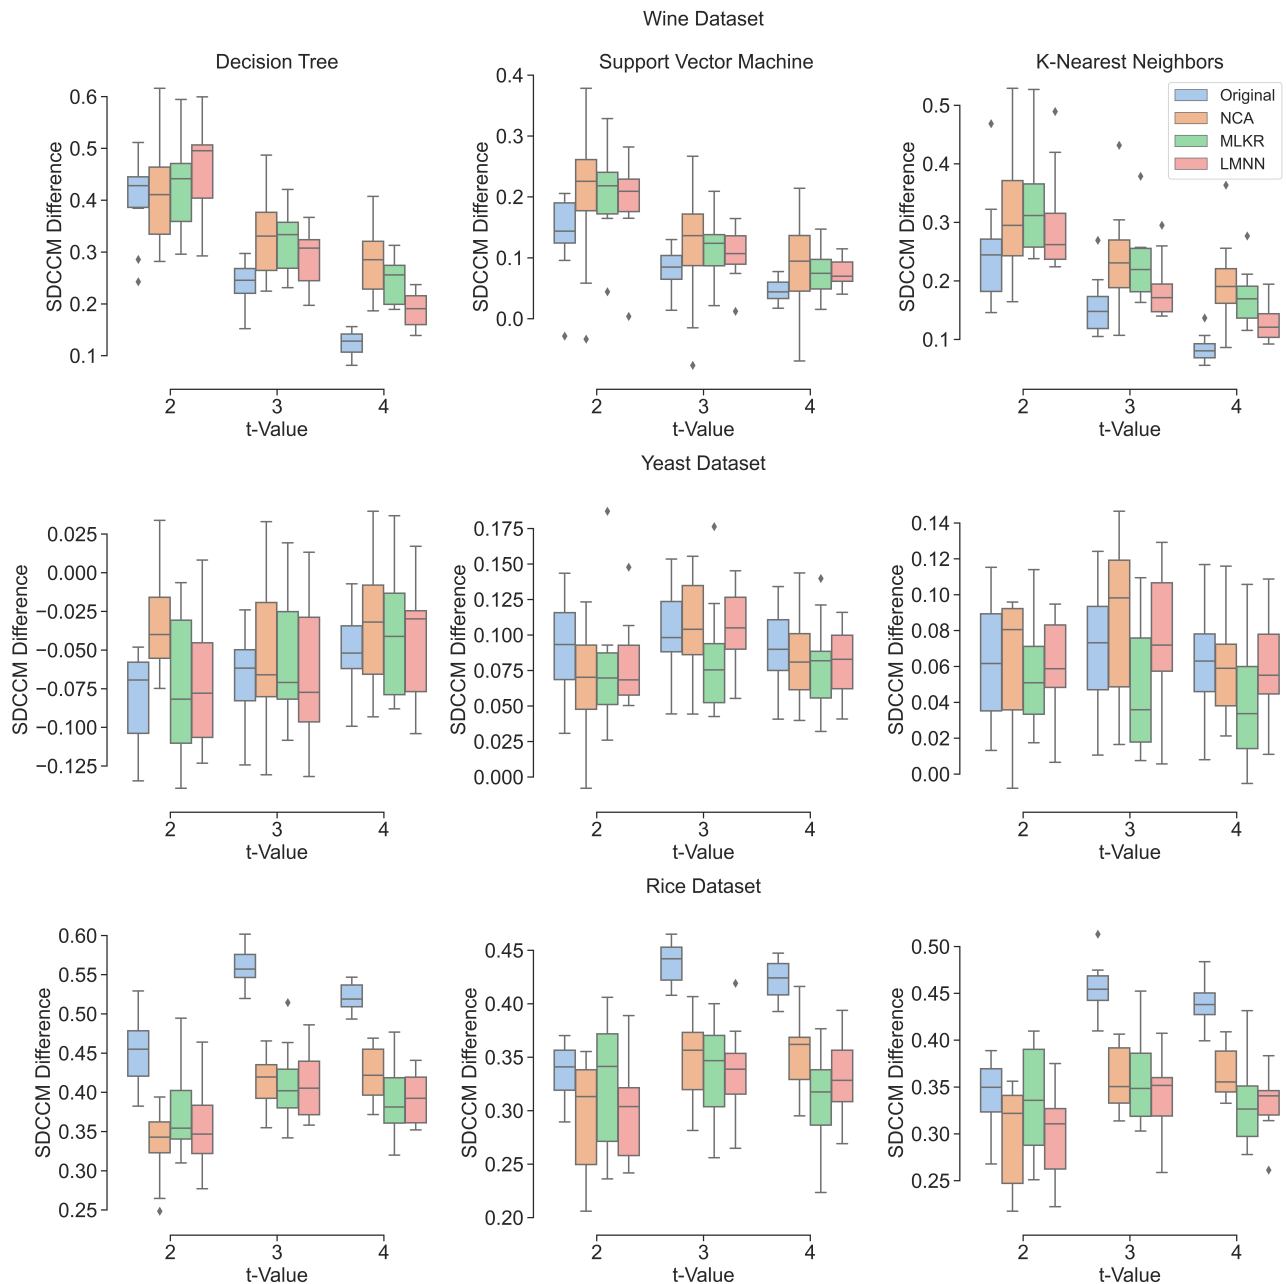
Fig. 1: Results by dataset and classifier for Wine, Yeast, and Rice.

smaller dependencies on t-values and classifiers, respectively. It should be noted, that across all datasets, classifiers, and t-values, only the Yeast dataset with the DT classifier shows near-zero difference between the SDCCMs for correctly and incorrectly classified test data. This provides evidence for claims in previous literature that SDCCMs can be used to anticipate error in machine learning [12]. While Cody et. al [13] demonstrated this correlation on MNIST image data, herein, we evidence the presence of a more general trend by using data from 6 different domains.

To better assess whether metric learning increases the ability

of SDCCMs to anticipate classification error, we conduct a rigorous statistical analysis using the paired t-test. The paired t-test assesses whether or not the pairwise differences between two sets of values varies from 0 [31]. If the test statistic is greater or less than 0 and the p-value is deemed significant, than the null hypothesis that there is no difference is deemed false. For our analysis, we treat test results with a p-value < 0.05 as significant. We pairwise compare the SDCCM differences of each metric learning space with those of the original space. We apply the paired t-test for each dataset and t-value since these are the primary dependencies identified in
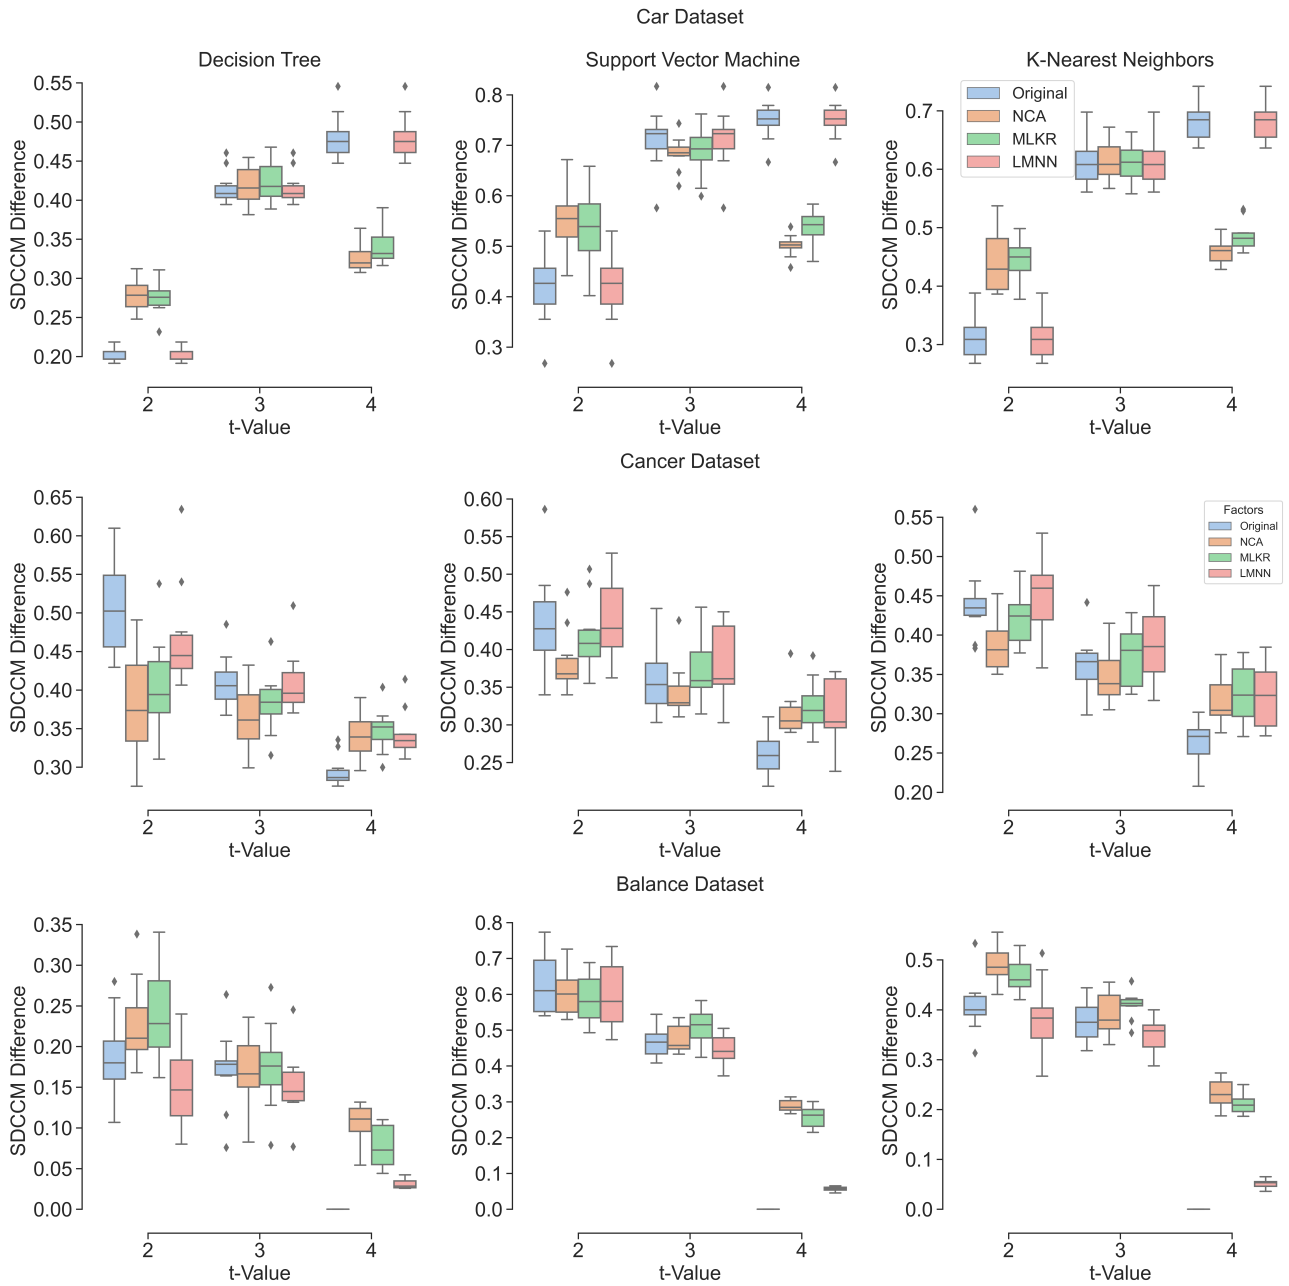
Fig. 2: Results by dataset and classifier for Car, Cancer, and Balance.

the boxplots.

In Table III, the results of the statistical analysis are shown with two significant figures. Since our hypothesis is that metric learning improves the ability of SDCCMs to anticipate classification error, as opposed to a specific metric learning method, the table identifies the metric learning method with the highest test statistic for each dataset and t-value. For the Car dataset at t-values of 2 and 3, the LMNN and original spaces were identical, and thus there was no pairwise difference. The results of Table III are summarized in Table IV. In $15/18$ or 83% of cases, there was a metric learning

space that was as good as or better than the original space. In $10/18$ or 56% of cases, there was a metric learning space that was better than the original space. In conclusion, these results provide evidence that metric learning improves the ability of SDCCMs to anticipate classification error.

VII. DISCUSSION

Results are discussed in detail in the previous section, however, the role of classification accuracy was not considered. Figure 3 shows the classification accuracy plotted against the SDCCM differences for the metric learning and original spaces. The points are from all datasets, t-values,

TABLE III
PAIRED T-TESTS FOR EACH DATASET AND t-VALUE

n = 30	t = 2				t = 3				t = 4			
Dataset	Statistic	p-Value	Result	Method	Statistic	p-Value	Result	Method	Statistic	p-Value	Result	Method
Wine	8.2	5.2e-9	Higher	LMNN	9.0	6.0e-10	Higher	MLKR	13	2.6e-13	Higher	LMNN
Rice	-2.5	2.1e-2	Lower	MLKR	-12	1.8e-12	Lower	MLKR	-13	2.7e-13	Lower	NCA
Yeast	1.1	2.9e-1	Same	NCA	2.2	3.5e-2	Higher	NCA	0.70	5.0e-1	Same	NCA
Car	18	1.2e-17	Higher	NCA	0	0	Same	LMNN	0	0	Same	LMNN
Cancer	-0.9	3.8e-1	Same	LMNN	2.3	2.9e-2	Higher	LMNN	13	2.5e-13	Higher	MLKR
Balance	2.7	1.1e-2	Higher	NCA	4.8	4.0e-5	Higher	MLKR	20	3.5e-18	Higher	LMNN

TABLE IV
PAIRED T-TEST SUMMARY

Dataset	t = 2	t = 3	t = 4	Totals
Wine	Higher	Higher	Higher	*Higher*
Rice	Lower	Lower	Lower	10
Yeast	Same	Higher	Same	*Same*
Car	Higher	Same	Same	5
Cancer	Same	Higher	Higher	*Lower*
Balance	Higher	Higher	Higher	3

and classifiers. The R^2 correlation between the accuracies and differences for all data (for each respective space) is displayed in the subplots. Also, regression lines are calculated for each dataset using linear least-squares regression. The R^2 correlations given by these regressions are displayed in the legend boxes of each subplot. The regression lines and their correlations are labeled by their dataset in the legend at the bottom of the figure.

Immediately, we can note that there are clear positive correlations between accuracy and SDCCM differences in all spaces. This means that the ability of SDCCM to anticipate classifier error increases with classification accuracy. Also, it appears that NCA and MLKR are able to transform points away from the upper-left quadrant, leading to higher correlations. These upper-left points indicate cases where SDCCM differences are uncorrelated to classification accuracy. This means that the SDCC calculations for correctly and incorrectly classified test data are similar regardless of classification accuracy.

It may seem undesirable to have points in the upper-left quadrant. But this intuition requires nuance. Consider the distinct goals of learning classifiers and learning latent spaces for coverage analysis. First, note that in this paper we are studying whether metric learning provides more meaningful spaces over which to conduct coverage analysis. Now, setting that motivation aside for a moment, consider that one has obtained a meaningful space—a space where no available method offers significant improvement in the ability to anticipate errors using SDCCMs. In such a space, the presence of points in the upper-left quadrant that indicate high performance but a low difference between SDCCMs calculated on correctly and incorrectly classified test data can be taken as a signal that classification errors are occurring at random, or at least without significant bias towards any particular region in the space of interactions. In contrast, the presence of points in

the upper-right quadrant indicate that there is such a bias; that is, that even though classification accuracy is high, there are still regions in the space of interactions where the classifier is predictably wrong.

Therefore, NCA and MLKR, by finding transformations which shift points away from the upper-left quadrant, are effectively finding new representations for the space of interactions to which the classifiers are predictably wrong or otherwise biased. From this reasoning, it seems that the goal of metric learning for coverage analysis should be to increase the correlation between model performance and SDCCM differences, while the goal of learning classifiers should be to decrease the correlation by shifting test points to the upper-left quadrant.

There are a few interesting trends in the dataset-specific correlations. Rice and Cancer, the two datasets where metric learning performed the worst based on visual inspection of Figures 1 and 2, consistently have the lowest correlations. Moreover, for Rice, the dataset where metric learning performed the worst based on the statistical analysis summarized in Table IV, the original space has a significantly higher correlation than any metric learning space. And for Cancer, the correlations are similar between the original and metric learning spaces. In contrast, Wine and Balance—the best performing metric learning methods based on Table IV—have correlations that are higher for the metric learning spaces than for the original space. From these observations, it seems that the larger the difference there is in like-for-like correlations between original and metric learning spaces, the more metric learning improves the ability of SDCCMs to anticipate error. This trend, however, needs more datasets for further investigation.

VIII. CONCLUSION

In this paper, we provide evidence that metric learning can improve the ability of SDCCMs to anticipate classification error. Importantly, the metric learning spaces were minimally tuned. This suggests that the demonstrated outperformance by metric learning is not over-engineered and can likely increase further with hyperparameter optimization. This promising evidence supports arguments for combinatorial coverage in identifying when out-of-distribution (OOD) data is likely to degrade model performance.

Additionally, our experiments support previous claims in the literature the SDCCMs correlate to the error of machine
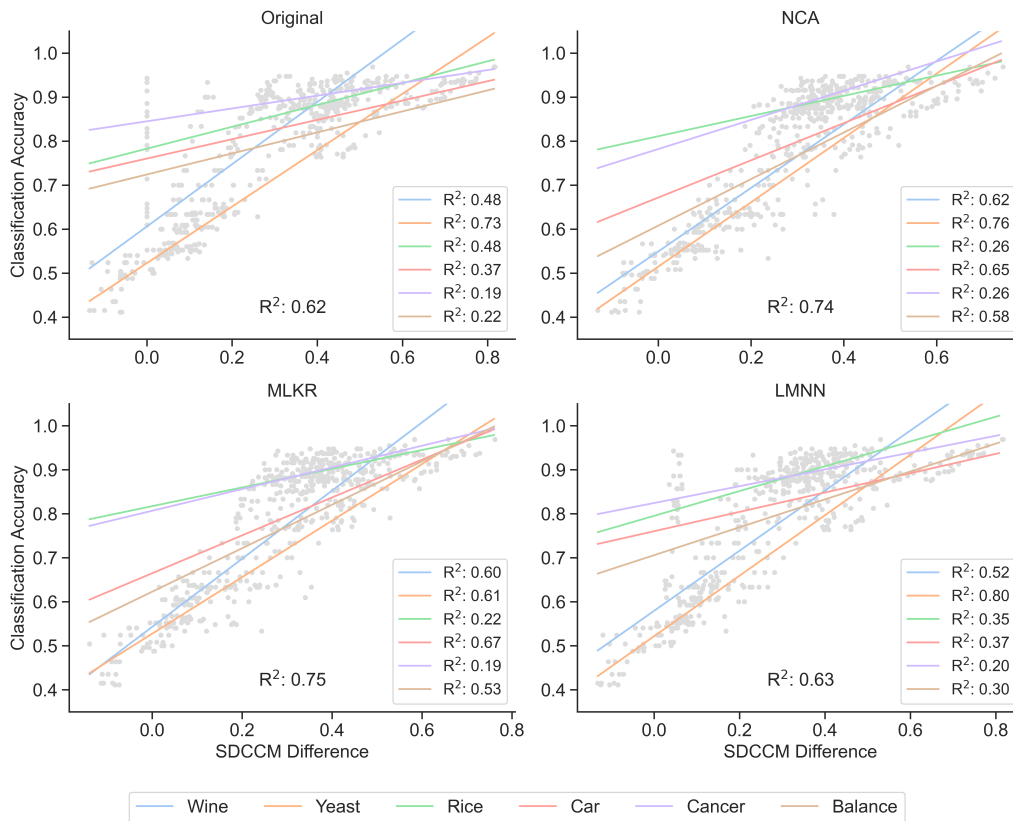
Fig. 3: Correlation analysis of classification accuracy and SDCCM differences.

learning classifiers. Despite past criticisms of particular applications of coverage concepts to neural networks, there is a growing body of literature supporting the use of combinatorial coverage concept in machine learning in general. As this marked trend continues, it is important to revisit the use of coverage concepts in learning algorithm development, not just auxiliary support systems like explainability, fairness, or error anticipation.

Metric learning offers a promising opportunity in this direction. Metric learning algorithms should be revisited in future work for their prospective use in coverage analysis. A particular concern in this area is metric learning methods that inherently learn discretized or readily discretizable spaces. Further, additional terms can be considered in the loss functions of metric learning algorithms such that the learned spaces are well-suited for particular kinds of coverage analysis, as opposed to, e.g., the traditional focus on being well-suited for classification and prediction problems. While the use of classical metric learning algorithms herein demonstrates an improvement in the ability to anticipate classification error, the results also indicate that significant dataset-dependence remains. Metric learning algorithms more tailored to coverage concepts could help further alleviate dataset-dependence in the utility of coverage analysis.

REFERENCES

[1] S. Arlot and A. Celisse, "A survey of cross-validation procedures for model selection," *Statistics Surveys*, vol. 4, pp. 40–79, 2010.
[2] P. Tripathy and K. Naik, *Software testing and quality assurance: theory and practice*. John Wiley & Sons, 2011.
[3] D. R. Kuhn, I. D. Mendoza, R. N. Kacker, and Y. Lei, "Combinatorial coverage measurement concepts and applications," in *2013 IEEE Sixth International Conference on Software Testing, Verification and Validation Workshops*. IEEE, 2013, pp. 352–361.
[4] D. R. Kuhn, D. R. Wallace, and A. M. Gallo, "Software fault interactions and implications for software testing," *IEEE Transactions on Software Engineering*, vol. 30, no. 6, pp. 418–421, 2004.
[5] K. Pei, Y. Cao, J. Yang, and S. Jana, "Deepxplore: Automated whitebox testing of deep learning systems," in *Proceedings of the 26th Symposium on Operating Systems Principles*, 2017, pp. 1–18.

[6] L. Ma, F. Juefei-Xu, F. Zhang, J. Sun, M. Xue, B. Li, C. Chen, T. Su, L. Li, Y. Liu *et al.*, "Deepgauge: Multi-granularity testing criteria for deep learning systems," in *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*, 2018, pp. 120–131.

[7] L. Ma, F. Juefei-Xu, M. Xue, B. Li, L. Li, Y. Liu, and J. Zhao, "Deepct: Tomographic combinatorial testing for deep learning systems," in *2019 IEEE 26th International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE, 2019, pp. 614–618.

[8] Z. Li, X. Ma, C. Xu, and C. Cao, "Structural coverage criteria for neural networks could be misleading," in *2019 IEEE/ACM 41st International Conference on Software Engineering: New Ideas and Emerging Results (ICSE-NIER)*. IEEE, 2019, pp. 89–92.

[9] D. R. Kuhn, R. N. Kacker, Y. Lei, and D. E. Simos, "Combinatorial methods for explainable AI," in *2020 IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW)*. IEEE, 2020, pp. 167–170.

[10] C. E. Tuncali, G. Fainekos, H. Ito, and J. Kapinski, "Simulation-based adversarial test generation for autonomous vehicles with machine learning components," in *2018 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2018, pp. 1555–1562.

[11] S. Ackerman, O. Raz, and M. Zalmanovici, "Freaai: Automated extraction of data slices to test machine learning models," in *International Workshop on Engineering Dependable and Secure Machine Learning Systems*. Springer, 2020, pp. 67–83.

[12] E. Lanus, L. J. Freeman, D. R. Kuhn, and R. N. Kacker, "Combinatorial testing metrics for machine learning," in *2021 IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW)*. IEEE, 2021, pp. 81–84.

[13] T. Cody, E. Lanus, D. D. Doyle, and L. Freeman, "Systematic training and testing for machine learning using combinatorial interaction testing," in *2022 IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW)*. IEEE, 2022, pp. 102–109.

[14] D. R. Kuhn, M. Raunak, C. Prado, V. C. Patil, and R. N. Kacker, "Combination frequency differencing for identifying design weaknesses in physical unclonable functions," in *2022 IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW)*. IEEE, 2022, pp. 110–117.

[15] J. Chandrasekaran, Y. Lei, R. Kacker, and D. R. Kuhn, "A combinatorial approach to explaining image classifiers," in *2021 IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW)*. IEEE, 2021, pp. 35–43.

[16] ——, "A combinatorial approach to testing deep neural network-based autonomous driving systems," in *2021 IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW)*. IEEE, 2021, pp. 57–66.

[17] A. R. Patel, J. Chandrasekaran, Y. Lei, R. N. Kacker, and D. R. Kuhn, "A combinatorial approach to fairness testing of machine learning models," in *2022 IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW)*. IEEE, 2022, pp. 94–101.

[18] T. Cody, J. Kauffman, J. Krometis, D. Sobien, and L. Freeman, "Combinatorial coverage framework for machine learning in multi-domain operations," in *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications IV*, vol. 12113. SPIE, 2022, pp. 467–471.

[19] S. Katragadda, T. Cody, P. A. Beling, and L. Freeman, "Active learning with combinatorial coverage," in *2022 21st IEEE International Conference on Machine Learning and Applications (ICMLA)*. IEEE, 2022, pp. 1100–1108.

[20] B. Kulis *et al.*, "Metric learning: A survey," *Foundations and Trends® in Machine Learning*, vol. 5, no. 4, pp. 287–364, 2013.

[21] S. Ramírez-Gallego, S. García, H. Mouriño-Talín, D. Martínez-Rego, V. Bolón-Canedo, A. Alonso-Betanzos, J. M. Benítez, and F. Herrera, "Data discretization: taxonomy and big data challenge," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 6, no. 1, pp. 5–21, 2016.

[22] M. Kaya and H. Ş. Bilge, "Deep metric learning: A survey," *Symmetry*, vol. 11, no. 9, p. 1066, 2019.

[23] K. Roth, T. Milbich, S. Sinha, P. Gupta, B. Ommer, and J. P. Cohen, "Revisiting training strategies and generalization performance in deep metric learning," in *International Conference on Machine Learning*. PMLR, 2020, pp. 8242–8252.

[24] K. Musgrave, S. Belongie, and S.-N. Lim, "A metric learning reality check," in *European Conference on Computer Vision*. Springer, 2020, pp. 681–699.

[25] J. Goldberger, G. E. Hinton, S. Roweis, and R. R. Salakhutdinov, "Neighbourhood components analysis," *Advances in Neural Information Processing Systems*, vol. 17, 2004.

[26] K. Q. Weinberger and G. Tesauro, "Metric learning for kernel regression," in *Artificial Intelligence and Statistics*. PMLR, 2007, pp. 612–619.

[27] K. Q. Weinberger and L. K. Saul, "Distance metric learning for large margin nearest neighbor classification." *Journal of machine learning research*, vol. 10, no. 2, 2009.

[28] A. Asuncion and D. Newman, "UCI machine learning repository," 2007.

[29] W. de Vazelhes, C. Carey, Y. Tang, N. Vauquier, and A. Bellet, "metric-learn: Metric Learning Algorithms in Python," *Journal of Machine Learning Research*, vol. 21, no. 138, pp. 1–6, 2020.

[30] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg *et al.*, "Scikit-learn: Machine learning in python," *the Journal of machine Learning research*, vol. 12, pp. 2825–2830, 2011.

[31] H. Hsu and P. A. Lachenbruch, "Paired t test," *Wiley StatsRef: Statistics Reference Online*, 2014.