Pairing Machine Learning and Clinical Psychology: How You Evaluate Predictive Performance
Matters

Ross Jacobucci¹, Andrew K. Littlefield², Alexander J. Millner³, Evan M. Kleiman⁴, & Douglas
Steinley⁵

¹ University of Notre Dame, Department of Psychology

² Texas Tech University, Department of Psychology

³ Harvard University, Department of Psychology

⁴ Rutgers, The State University of New Jersey, Department of Psychology

⁵ University of Missouri, Department of Psychology

Please address your correspondences to Ross Jacobucci, Department of Psychology, University of Notre
Dame, Notre Dame, IN 46556.

A preprint and all materials used for the empirical example and simulation is made publicly available at

<https://osf.io/73pnk/>

Abstract

Machine learning is being utilized at an increasing rate in clinical psychology. Applying machine learning comes with a number of challenges, both in deciding which algorithms to test, and how to evaluate the predictive performance. We focus on this last component, demonstrating across both a simulation and empirical example that the method researchers choose to evaluate prediction with machine learning can have large consequences for the substantive conclusions. More specifically, we demonstrate that one method of evaluation that has been used repeatedly in clinical research, the optimism corrected bootstrap, can result in extremely biased results when paired with specific machine learning algorithms. We conclude with providing recommendations for researchers and a discussion of additional best practices.

Keywords: machine learning, data mining, prediction, clinical psychology

Introduction

There is increasing support behind the idea that psychiatric phenomena are highly complex and therefore cannot be explained by a small number of variables with linear relationships (Kendler, 2005; Kendler, 2019; Franklin, 2019 Mitchell, 2015). Perhaps one of the most complex phenomena is suicide, an outcome that is difficult to predict and appears to have many causes. One approach to dealing with such complexity is to use advanced statistical approaches, such as machine learning (similar terms include data mining, artificial intelligence, or statistical learning; for an overview, see Dwyer, Falkai, & Koutsouleris, 2018), with large numbers of variables to predict an important outcome and, in some cases, to determine which variables contribute the most information for such predictions. Indeed, machine learning algorithms¹ have been utilized to increase predictive accuracy and identify novel indicators within suicide research (see Burke, Ammerman, & Jacobucci, 2018), including influential work published in this journal focused on predicting suicidal outcomes (e.g., Walsh, Ribeiro, & Franklin, 2017; Ribeiro, Huang, Fox, Walsh, & Linthicum, 2019).

A primary rationale behind the use of machine learning approaches is to account for complex interactions or nonlinear effects among psychological variables. Traditional approaches, such as logistic regression, cannot as feasibly account for such effects and therefore are assumed to have poorer predictive capability compared with these advanced modeling approaches, particularly when studying complex phenomena. For example, when a recent meta-analysis examining the prediction of suicidal thoughts and behaviors over the past fifty years of research found only small predictive effects (Franklin et al., 2017), the authors argued that suicide prediction required “a shift in focus from risk factors to machine learning-based risk algorithms” (p. 187).

The argument for machine learning algorithms has been supported by multiple studies where machine learning approaches outperformed more traditional methods in predicting suicidal behaviors. Walsh et al. (2017) found that a machine learning algorithm (random forests) outperformed standard logistic regression models (as indexed by the area under the receiver operator characteristic area under the curve [AUC]) when used to predict risk of suicide attempts (i.e., random forest AUC were .80 - .84 versus .66 - .68 for multiple logistic regression models; see Walsh et al., 2017, for more details). Other studies compared random forests and logistic regression when predicting suicidal thoughts and behaviors (Ribeiro et al., 2019), non-suicidal self-injury over short time windows (within 2-4 weeks; Fox et. al., 2019) and when identifying variables that distinguish those that think about suicide from those that attempt suicide (Huang, Riberio, and Franklin, 2019). In each of these studies, the authors found that random forests outperformed single and multiple logistic regression, concluding that results supported the idea that the outcome of interest was too complex to be explained through simple (e.g., linear) modeling approaches.

However, some studies have failed to find an advantage for machine learning approaches to predict suicide outcomes, compared to logistic regression. For example, a recent study with a large sample ($N = 3,649$) that sought to determine a suicide risk index among Army soldiers using prospective data found, “no advantage to using a machine learning model that allowed for interactions over a simple logistic regression model” (p. 6; see Zuromski et al., 2019). These findings within the field of suicidality are consistent with results from a large systematic review that concluded, across a range of medical conditions with ostensibly complex etiologies (e.g., oncology), there was no performance benefit of machine learning over logistic regression for clinical prediction models (see Christodoulou et al., 2019). Thus, before fully embracing the

theoretical conclusions drawn from machine learning studies (e.g., that suicide risk is best modeled as a complex process that requires machine learning), as well as using these models for risk assessment, it is important to understand the variability in study outcomes with regard to the comparative performance of machine learning models.

Several factors could account for the discrepancy in predictive performance across studies using machine learning. These include sample size, sample characteristics, variables used (in terms of both the number of variables considered as well as measurement quality; see Jacobucci & Grimm, *in press*), the algorithms assessed, and the form of validation used to assess the performance. Our point of emphasis in this paper is this last point, precisely because it can have a large impact, as we will show, and it is within the researcher's control (i.e., falling under the umbrella of researcher degrees of freedom; Simmons, Nelson, & Simonsohn, 2011). A crucial part of evaluating model performance concerns the generalizability of a given algorithm to new data. Dwyer, Falkai, and Koutsouleries (2018) outline a hierarchy of generalizability assessment, strongly advocating for the use of external, multi-site, or prospective validation of the derived model. However, most researchers do not have access to an alternate sample to perform external validation, which is made increasingly difficult given the data hungry nature of machine learning algorithms (van der Ploeg, Austin, & Steyerberg, 2014). Flexible machine learning algorithms can demonstrate strong prediction results by capitalizing on patterns within a given data set that do not generalize to new data (i.e., these models can grossly overfit the data; see Domingos, 2012, Yarkoni & Westfall, 2017, for further detail).

Given the use of a single dataset, there are a multitude of approaches researchers use to assess, and thus prevent, overfitting, with the most common being cross-validation (CV) and bootstrap sampling. Both strategies iterate through the data using different data subsets as the

training set and another subset to test the performance of an algorithm. That is, these approaches attempt to avoid overfitting not by testing in a new dataset an algorithm developed in one dataset, but rather testing an algorithm using a subset of the same dataset in which the algorithm was developed. There are many variations on both CV and bootstrap sampling, such as k -fold and repeated k -fold CV, as well as the bootstrap and the optimism-corrected bootstrap, to name a few. The studies highlighted above demonstrate the variability in the use of internal validations methods, with some papers using k -fold CV and some using the optimism-corrected bootstrap. Consequently, to determine why some papers report that machine learning algorithms outperform traditional statistical models, and others come to opposite conclusions, examining whether all strategies guard equally against overfitting is necessary. Detailed descriptions of these resampling approaches are provided below.

Cross-Validation

One can perform numerous variants of CV, with the most common being k -fold CV. This involves repeatedly splitting the sample into k partitions (commonly 5 or 10), running (often referred to as “training”) the model on $k-1$ combined partitions of the data, and then testing the model on the single partition that was held out. Note that this latter part in testing the model means fixing the model estimates to the values from the training sets to create predictions on the holdout (e.g., test) set, not re-estimating the model. We depict the process of k -fold CV (specifically, five-fold CV) in Figure 1. In this example, we can assign each observation to one of five partitions, to which we repeat the process of training the model on a combination of four partitions k times, and test (treat model parameters as fixed) on the remaining partition each time.

To get an overall assessment of model fit, the performance metrics (such as the AUC) can then be averaged across the k iterations. Given that k -fold CV involves only one initial random

partition of the data, repeating the process (i.e., 10-100) of partitioning the data is often recommended (Borra & Di Ciaccio, 2010; Krstajic, Buturovic, Leahy, & Thomas, 2014; Kuhn & Johnson, 2013), though this can add a computational burden.

Bootstrap Sampling

Bootstrap sampling (Efron & Tibshirani, 1986) involves sampling with replacement from the original sample to create a bootstrap sample. This new sample of size N will contain some observations multiple times, with approximately 36.8% of the observations not being selected for the bootstrap sample (“out of bag” sample). This holdout set can then be utilized as a test set in a similar way as with cross-validation, while training on the bootstrap sample. Repeating this process multiple times (e.g., 20-200) allows researchers to derive reasonably accurate estimates of predictive performance.

A number of variations exist to bootstrap sampling, many involving balancing the predictive performance on the bootstrap sample (the training set; apparent error rate) and the observations held out. Of particular interest, given its use in a number of applications of machine learning to clinical psychology research (e.g., Austin, Lee, Steyerberg, & Tu, 2012; Esmaily, Tayefi, Doosti, Nezami, & Amirabadizadeh, 2018; Huang, Ribeiro, & Franklin, 2019; Matsuki, Kuperman, & Van Dyke, 2016; Ribeiro, Huang, Fox, Walsh, & Linthicum, 2019; Walsh, Ribeiro, & Franklin, 2017), is the optimism corrected bootstrap. This algorithm proceeds by taking the following steps:

1. Create a model on the original sample, calculating performance
2. Construct models on a sequence of bootstrap samples

3. Test the model created on each bootstrap model back on the original sample (test performance)
4. Calculate *optimism* by subtracting test performance from bootstrap performance, averaging these values to create a mean optimism estimate
5. Subtract the mean optimism estimate from the performance as assessed on the original sample (step 1)

As previously mentioned, this resampling method has been used extensively in practice and is often cited as a relatively conservative approach (e.g., Walsh et al., 2017) based on performance when using simpler approaches such as logistic regression (e.g., Hastie, Tibshirani, & Friedman, 2001). However, a number of studies have found suboptimal performance (John, 2018; Tantithamthavorn, McIntosh, Hassan, & Matsumoto, 2016), most notably that this form of bootstrap can be biased positively. More specifically, Tantithamthavorn et al. (2001) found coupling this approach with random forests caused gross overestimates of model performance, suggesting validation approaches should be examined in tandem with specific modeling approaches. Thus, we first detail three of the most commonly applied algorithms; elastic net, boosting, and random forests.

Algorithms

Prior to describing these algorithms, we first define the term *tuning parameter* (i.e., hyperparameter). Most machine learning algorithms have a multitude of settings that control various aspects of their flexibility and thus performance. These settings are often termed tuning parameters, as they are aspects of the algorithms that are not set but are necessary to calibrate (tune) to determine which specific value produces the best performance. The number of tuning

parameters can vary across algorithms, along with the number of distinct values that are worth testing. We further define this in the context of each algorithm.

Elastic net regression (Zou & Hastie, 2005) can be viewed as a generalization of linear regression that adds a penalty to the regression coefficients, thus performing variable selection (the lasso penalty component, where lasso stands for least absolute shrinkage and selection operator), while also handling multicollinearity (the ridge penalty component). The elastic net builds upon linear regression, using the same residual sum of squares equation as a basis, while including two tuning parameters that control both the ridge and lasso penalty contributions. The relative contributions of both the ridge and lasso penalties is controlled by a mixing parameter, often termed *alpha*. In lasso regression, when the penalty becomes large enough, a subset or all regression coefficients will be set to zero, thus performing variable selection. The ridge penalty does not result in variable selection but is an important component to include as it better controls for collinearity among the predictors.

Both random forests (Breiman, 2001) and boosting (Friedman, 2002; Freund, & Schapire, 1997) are ensemble methods (i.e., methods that combine several machine learning methods together), building a large number of individual decision trees, aggregating predictions across these hundreds or thousands of trees. Decision trees (Morgan & Sonquist, 1963; Breiman, Friedman, Olshen, & Stone, 1984) recursively partition the predictor space to identify homogenous subgroups that have similar responses on an outcome. The decision tree algorithm can efficiently identify interactions and nonlinear effects when the number of predictors is large. However, the interpretability of decision trees comes at the expense of stability and predictive performance, thus spawning the creation of more complex algorithms.

Random forests (Breiman, 2001) can be thought of as an aggregation of many individual trees (e.g., 100), creating each tree with a bootstrap (or subset) of the sample and a subset of the predictors. After creating a large number of trees, predictions are made for each individual by averaging predictions across trees, which is known as bootstrap aggregation (bagging; Breiman, 1996b). This results in a complex relationship between predictors and the response that can fit highly nonlinear relationship. The main tuning parameters to random forests are the fraction of predictors selected for creating each tree (typically referred to as *mtry*), as well as the number of trees.

Boosting (Freund & Schapire, 1997; Friedman, 2002) also creates many trees, but creates each tree in a sequential manner, with each tree fit to the residuals that results from the previously created set of trees. An additional distinction with random forests is that boosting typically involves creating smaller trees, with the tree depth being commensurate to the expected degree of interactions between variables (Elith, Leathwick, & Hastie, 2008). Additionally, regularization is incorporated into the algorithm, which limits the influence of any single tree, producing a slower learning process (thus typically requiring more trees than random forests). This is akin to the regularization component in elastic net regression, which limits the influence of any single variable. Taken together, boosting incorporates a larger number of tuning parameters, including tree depth, penalty value (shrinkage or regularization), and number of trees.

Our purpose in this paper is to highlight a problematic form of internal validation, namely the optimism corrected bootstrap, by comparing the performance of this validation technique (using both simulated and actual data) to variants of CV and the standard bootstrap. Importantly,

given that the performance of internal validation methods may vary across types of algorithms, we compare the validation techniques across elastic net, boosting, and random forest algorithms.

Simulation

To understand how the form of resampling chosen can influence conclusions with respect to predictive performance, we performed a small-scale simulation study. As influencing conditions, we hypothesized that both the sample size and number of predictors would result in biased performance metrics when the sample size was small and the number of predictors large. As a base model, we simulated a linear regression model with a sample size of 20,000 and an outcome Y resulting from two standard normal X variables with coefficients of 0.3 and standard normal error. This continuous Y variable was then dichotomized (based on Z -score of $-.5$, resulting in approximately 32% 1's) to mimic the common use of categorical outcomes in clinical psychology research.

Treating the sample of 20,000 as the population, we repeatedly subsampled from this to sample sizes of 100, 200, 500, 1,000, and 2,000 (training sample), using the remaining samples as a “true” holdout, using the performance on this sample as an approximation to ground truth prediction performance. An additional simulation condition was the number of added noise variables, simulated standard normal variables with no association with Y , varied across 5, 20, 50, and 100 variables. We believe these simulation conditions to be representative of typical machine learning studies in clinical research².

In comparison to the elastic net, both boosting and random forests are much more flexible in the functional forms that they can approximate, thus leading to a higher propensity to overfit. We expect more bias for both boosting and random forests, particularly when the sample size is small or the number of noise variables is large, if any of the resampling methods do not

adequately measure overfitting with respect to holdout performance. To evaluate holdout performance, algorithms were trained on the training sample, a final model was selected according to the tuning parameter that evidenced the best resampling performance, then these model parameters were treated as fixed and used to create predictions on the holdout sample. We compared this ground truth performance metric to the performance as assessed with 10-fold CV, repeated 10-fold CV (with 10 repeats), the bootstrap, and the optimism corrected bootstrap.

As a follow up to this main set of simulations, we evaluated the performance of removing one of the two “true” predictors prior to running random forests to determine whether each form of resampling indicates a substantial drop in AUC. This strategy is in line with how previous research has attempted to identify whether there are a small number of important predictors that contribute to modeling the outcome, or whether modeling the outcome requires a large set of predictors, each contributing a relatively small amount (complexity; Ribeiro, Huang, Fox, Walsh, & Linthicum, 2019). For this set of simulations, we only compared 10-fold CV and the Optimism Corrected Bootstrap while varying the sample size and number of noise predictors.

All analysis was performed in the R statistical environment (R Core Team, 2019) with `train()` from the `caret` package (Kuhn, 2008; Kuhn & Johnson, 2013). `train()` acts as a wrapper for a large number of algorithms and resampling methods. We were able to test each form of resampling, as well as use the elastic net, random forests, and boosting. The `caret` package also allows for the specification of the number of distinct values used for each tuning parameter, while generating defaults for each of the tuning parameters. For this, we detail the number of tuning parameters, while using defaults for each. More specifically, we used the following implementation of each algorithm. We used the `glmnet` package (Friedman, Hastie, & Tibshirani, 2010) implementation of the elastic net, testing 30 distinct values (using default

values) of both alpha (the mixing) and the penalty, resulting in 900 different models. For random forests, we used the `randomForest` package (Liaw, & Wiener, 2002) implementation, and three values of `mtry`³ and 500 trees. For boosting, we used the `gbm` package (Greenwell, Boehmke, Cunningham, & GBM Developers, 2019) implementation, with 3 values of tuning parameters for number of trees, tree depth, and shrinkage (regularization penalty), resulting in a total of 27 distinct models.

Simulation Results

Results from the first set of simulations are displayed in Figures 2-7. In Figures 2 and 3 for the elastic net, we can see that each form of resampling results in similar conclusions regarding the performance across sample size and number of noise predictors. The only deviations occur with the optimism corrected bootstrap when the sample size is small (100 or 200) or the number of noise predictors is large (50 or 100). In contrast are the results from using boosting and random forests (Figures 4-7), with upward biased estimates of AUC when using the optimism corrected bootstrap regardless of the number of predictors or sample size. For both random forests and boosting, this is particularly exacerbated when the sample size is small, and the number of noise predictors is large.

In collapsing across simulation conditions, we can calculate mean bias for each of the validation methods. The degree of bias to each method can be quantified by examining the deviations from holdout performance (“Holdout” performance is as close to the truth as possible, as it is evaluated on new data). While the forms of cross-validation were only slightly positively biased (0-0.02 for 10-fold CV), and the standard bootstrap was essentially unbiased, the optimism bootstrap demonstrated much higher degrees of bias, with values of 0.05 (meaning the AUC was on average 0.05 higher than assessed in the holdout) for the elastic net, 0.22 for

boosting, and 0.33 for random forests. The most biased results occurred for random forests combined with the optimism corrected bootstrap, with estimates of AUC nearing .95 across simulation conditions, when holdout AUC was estimated as ranging from approximately .60-.65. Of note is the larger amount of variability to the AUC estimates when the sample size is smaller, regardless of resampling type, meaning that although the average estimates are close to holdout performance, the substantial variability will lead to upwardly biased conclusion in a subset of cases.

For the second set of simulation results examining the removal of one of the “true” predictors, we see quite similar results in Figure 8 to the first set of simulations, with the optimism corrected bootstrap demonstrating extremely positively biased performance, where there are small differences in performance between one and two predictors. This occurs despite the removal of a contributing factor for explaining the outcome. In contrast, 10-fold CV shows a substantial decrement in performance between two and one predictor (noise variables were included) models, particularly at larger sample sizes. These results demonstrate that pairing random forests and the optimism corrected bootstrap can also bias the assessment of which predictors are important.

Empirical Example

Resampling techniques and modeling approaches used in the simulation analyses were compared using the baseline data from National Comorbidity Survey (NCS; Kessler et al., 1994). A prior analysis of this dataset used random forests, paired with the optimism corrected bootstrap, to model suicide attempter vs. suicidal ideation status (Franklin, 2019, refers to work conducted by Huang, Ribeiro, and Franklin available at <https://psyarxiv.com/8tuqg/>, which demonstrated that ideators and attempters within the NCS dataset can be much more accurately

distinguished using random forests [AUC = .89] compared to logistic regression [AUC = .71]). For pedagogical purposes, models were based on the 51 independent variables (no manual inclusion of interactions or quadratic effects) described in Huang et al. (2019). These predictors consisted of demographic variables (e.g., age, sex), psychiatric diagnoses (e.g., bipolar disorder diagnosis, generalized anxiety disorder diagnosis), negative life events (e.g., childhood neglect), personality (i.e., neuroticism), and a history of suicidal behaviors (e.g., suicide ideation age of onset; see Huang et al., 2019, for more details).

The total sample size for our analyses was 885 (563 ideators and 322 attempters). All machine learning analyses were conducted using the `caret` package (Kuhn, 2008), using the same R packages as used in the simulation analyses (i.e., `glmnet`, `randomForest`, and `gbm` for the respective elastic net, random forest, and gradient boosting machine analyses). In terms of tuning parameters, the elastic net included seven values (ranging from 0 to 1) for alpha, the mixing parameter, and 40 values (ranging from .01 to .20) of lambda, the penalty parameter. For random forest, `mtry` (the fraction of predictors selected for creating each tree) was set to seven with 500 trees. Gradient boosting machine tuning parameters considered three values for interaction depth (1, 2, 3), three values for the number of trees (100, 500, 2000), three values for shrinkage (.1, .01, .001), and the number of minimum observations in a node (`n.minobsinnode`) was held constant at 20. The “`optimism_boot`” option was used to calculate the optimism corrected bootstrap and the “`boot`” option was used to calculate a standard bootstrap (with 100 bootstrap draws for each approach) estimates. The “`cv`” option was used to calculate the 10-fold CV and the “`repeatedcv`” option was used to calculate 10-fold CV with 10 repeat estimates. Outcomes were optimized on AUC using the one standard error rule (i.e., choosing the most parsimonious model whose error is no more than one standard error above the error of the best

model; see Hastie, Tibshirani, & Friedman, 2009). We also included a linear logistic regression model as a point of comparison.

Empirical Example Results

AUCs across models and resampling techniques are shown in Table 1. Of the 16 AUC estimates across models and resampling strategies, 14 were quite similar (AUCs ranged from .77-.79), including AUCs from the logistic regression model. Consistent with the simulation findings, the AUC was markedly higher for the optimism corrected bootstrap approach when paired with gradient boosting machine (AUC = .93) and random forest (AUC = .97). These findings are consistent with the simulation results that demonstrated that the optimism corrected bootstrap results in inflated estimates of model performance when paired with gradient boosting machine and random forest. When using other resampling methods, the machine learning models performed generally as well as a traditional linear model.

Discussion

In this paper, we have given a brief overview of machine learning algorithms and internal validation strategies for assessing prediction performance. Ideally, researchers have access to an external dataset that can be used to accurately assess the performance of prediction models. However, for the majority of research, external datasets do not exist for a variety of reasons. Across both the simulation and empirical example, we detailed the biased performance of one internal validation method, the optimism corrected bootstrap, which has been used repeatedly in clinical research. Although there are a multitude of factors that can play a part in determining the performance of applying machine learning in clinical research, choice of internal validation strategy is directly in a researcher's control. We demonstrated that this choice can have large implications, which can then lead to inaccurate clinical conclusions.

In our simulation, the comparable performance of the optimism corrected bootstrap and other resampling methods when using regression, while appreciably worse performance (i.e., positively biased) of the optimism corrected bootstrap with random forests is in line with the findings of Tantithamthavorn, McIntosh, Hassan, and Matsumoto (2016). Additionally, the higher variability to the estimates of AUC across the three algorithms is in line with recommendations for the use of larger sample sizes when using machine learning (van der Ploeg, Austin, & Steyerberg, 2014). Given that the optimism corrected bootstrap has been used as the sole method of validation across several high impact studies within clinical psychology (including two articles published in this journal), there are several implications of our current findings.

Practically, there needs to be greater focus on the type of validation approach used by researchers in the field. This recommendation echoes the suggestions made in Christodoulou et al. (2019) regarding the utilization of machine learning for clinical prediction within the medical field. More specifically, these authors noted that model validation procedures are often unsound or not well reported; indeed, the authors found studies with these characteristics were more likely to indicate enhanced machine learning performance compared to traditional approaches. In light of these observations, Christodoulou et al. (2019) recommended a greater focus on validation procedures and assessments of model performance when comparing machine learning to traditional methods. Perhaps reassuringly, our current results indicate relatively good and comparable performance across k -fold and traditional bootstrapping approaches. Regardless, it may behoove the field if greater emphasis is placed on validation approaches (with perhaps the need of sensitivity analysis to determine the robustness of a given finding) before concluding a given piece of work convincingly demonstrates machine learning algorithms are superior to

traditional, less opaque modeling approaches. Further, it should not be assumed that a validation approach that has been shown to be valid for one type of algorithm is valid for other types of algorithms, as different machine learning algorithms have varying propensities to overfit the data.

The current work also has theoretical implications. The apparent marked superiority of machine learning approaches over traditional algorithms (when using optimism corrected bootstrap as the sole validation approach) has been cited as evidence that clinically relevant phenomenon (such as those related to suicidality) are quite complex, requiring large numbers of predictors and nonlinear relationships. Relatedly, if random forests or boosting evidence superior prediction performance, it provides rationale for then producing and interpreting metrics or graphics that assess nonlinearity or interaction effects, such as partial dependence plots. These visual depictions of relationships between one or two predictors and the outcome often depict elaborate, nonlinear relationships even when there is little evidence for their existence. Biased performance metrics can then lead to inaccurate conclusions regarding the types of relationships inherent in the data.

Importantly, it could be the case that prior work based on the optimism corrected bootstrap would also demonstrate markedly enhanced performance of machine learning when using less biased validation methods, which would maintain the support of complex relations between independent variables and outcomes. However, this remains to be demonstrated, although doing so would be quite easy as it would involve only a re-analysis of existing datasets. The possibility that the utility of certain machine learning algorithms to model complex relations among variables has been overstated (at least in some datasets) certainly raises several critical questions. For example, if reanalysis of prior work suggests a general linear model performs as

well as an approach that allows for complex interactions, does that provide evidence for the notion that certain psychological phenomenon (e.g., differences between suicide ideators and suicide attempters) are relatively simple rather than complex? To what extent will a shift in focus from risk factors to machine learning-based risk algorithms reliably advance the field of clinical psychology, at least when paired with traditional variables and research designs?

Despite the questions the current work raises, our purpose in writing this paper is not to dissuade researchers from using machine learning to predict clinical outcomes. Instead, our goal was to highlight the importance in how one evaluates the performance of machine learning algorithms, detailing vastly discrepant results across various forms of internal validation. We specifically show that one's choice of internal validation method can have large implications. Given that psychology is facing a replication crisis (Open Science Collaboration, 2015), our goal here is to prevent the same thing from happening in the use of machine learning in clinical psychology. As such, we recommend for researchers to use one of the internal validation methods (when external data is not available) that has demonstrated accurate performance across various machine learning algorithms. If a researcher is uncertain about the veracity of the results, we recommend using multiple types of resampling (e.g., both k-fold cross-validation and traditional bootstrap sampling).

References

- Austin, P. C., Lee, D. S., Steyerberg, E. W., & Tu, J. V. (2012). Regression trees for predicting mortality in patients with cardiovascular disease: What improvement is achieved by using ensemble-based methods?. *Biometrical Journal*, 54(5), 657-673.
- Borra, S., & Di Ciaccio, A. (2010). Measuring the prediction error. A comparison of cross-validation, bootstrap and covariance penalty methods. *Computational Statistics & Data Analysis*, 54(12), 2976-2989.
- Breiman, L. (2001). Random forests. *Machine learning*, 45(1), 5-32.
- Breiman, L., Friedman, J., Olshen, R., Stone, C., 1984. Classification and Regression Trees. Wadsworth International Group, Belmont, CA.
- Burke, T. A., Ammerman, B. A., & Jacobucci, R. (2018). The use of machine learning in the study of suicidal and non-suicidal self-injurious thoughts and behaviors: A systematic review. *Journal of Affective Disorders*, 245, 869-884.
- Buuren, S. V., & Groothuis-Oudshoorn, K. (2010). mice: Multivariate imputation by chained equations in R. *Journal of Statistical Software*, 45, 1-68.
- Christodoulou, E., Jie, M. A., Collins, G. S., Steyerberg, E. W., Verbakel, J. Y., & van Calster, B. (2019). A systematic review shows no performance benefit of machine learning over logistic regression for clinical prediction models. *Journal of Clinical Epidemiology*, 110, 12-22.
- Davis, J., & Goadrich, M. (2006, June). The relationship between Precision-Recall and ROC curves. In *Proceedings of the 23rd international conference on Machine learning* (pp. 233-240). ACM.

- Domingos, P. M. (2012). A few useful things to know about machine learning. *Communications of the ACM*, 55(10), 78-87.
- Dwyer, D. B., Falkai, P., & Koutsouleris, N. (2018). Machine learning approaches for clinical psychology and psychiatry. *Annual Review of Clinical Psychology*, 14, 91-118.
- Efron, B. (1983). Estimating the error rate of a prediction rule: improvement on cross-validation. *Journal of the American Statistical Association*, 78(382), 316-331.
- Efron, B., & Tibshirani, R. (1997). Improvements on cross-validation: the 632+ bootstrap method. *Journal of the American Statistical Association*, 92(438), 548-560.
- Elith, J., Leathwick, J. R., & Hastie, T. (2008). A working guide to boosted regression trees. *Journal of Animal Ecology*, 77(4), 802-813.
- Esmaily, H., Tayefi, M., Doosti, H., Nezami, H., & Amirabadizadeh, A. (2018). A comparison between decision tree and random forest in determining the risk factors associated with type 2 diabetes. *Journal of Research in Health Sciences*, 18, e00412.
- Fawcett, T. (2004). ROC graphs: Notes and practical considerations for researchers. *Machine Learning*, 31(1), 1-38.
- Fox, K. R., Huang, X., Linthicum, K. P., Wang, S. B., Franklin, J. C., & Ribeiro, J. D. (2019). Model complexity improves the prediction of nonsuicidal self-injury. *Journal of Consulting and Clinical Psychology*, 87, 684-692.
- Franklin, J. C. (2019). Psychological primitives can make sense of biopsychosocial factor complexity in psychopathology. *BMC Medicine*, 17(1), 187.
- Freund, Y., & Schapire, R. E. (1997). A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55, 119-139.

- Friedman, J. H. (2002). Stochastic gradient boosting. *Computational Statistics & Data Analysis*, 38, 367-378.
- Friedman, J., Hastie, T., & Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1), 1.
- Greenwell, B., Boehmke, B., Cunningham, J., & GBM Developers (2019). gbm: Generalized Boosted Regression Models. R package version 2.1.5. <https://CRAN.R-project.org/package=gbm>
- Harrell, F. (2015). Regression modeling strategies: with applications to linear models, logistic and ordinal regression, and survival analysis. Springer, New York.
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). The elements of statistical learning. Springer, New York.
- Huang, X., Ribeiro, J. D., & Franklin, J. C. (2019, August 29). The Differences between Suicide Ideators and Suicide Attempters: Simple, Complicated, or Complex?.
<https://doi.org/10.31234/osf.io/8tuqg>
- Jacobucci, R., & Grimm, K. (in press). Machine learning and psychological research: The unexplored effect of measurement. *Perspectives on Psychological Science*.
- John, C. (2018, December 25). Optimism corrected bootstrapping: a problematic method. [Blog post]. Retrieved from <https://intobioinformatics.wordpress.com/2018/12/25/optimism-corrected-bootstrapping-a-problematic-method/>
- Karstoft, K. I., Galatzer-Levy, I. R., Statnikov, A., Li, Z., & Shalev, A. Y. (2015). Bridging a translational gap: using machine learning to improve the prediction of PTSD. *BMC Psychiatry*, 15(1), 30.

- Kendler, K. S. (2005). Toward a Philosophical Structure for Psychiatry. *American Journal of Psychiatry*, 162(3), 433–440. <https://doi.org/10.1176/appi.ajp.162.3.433>
- Kendler, K. S. (2019). From Many to One to Many—The Search for Causes of Psychiatric Illness. *JAMA Psychiatry*. <https://doi.org/10.1001/jamapsychiatry.2019.1200>
- Krstajic, D., Buturovic, L. J., Leahy, D. E., & Thomas, S. (2014). Cross-validation pitfalls when selecting and assessing regression and classification models. *Journal of Cheminformatics*, 6, 1-15.
- Kessler, R. C., McGonagle, K. A., Zhao, S., Nelson, C. B., Hughes, M., Eshleman, S., Wittchen, H.-U., & Kendler, K. S. (1994). Lifetime and 12-Month Prevalence of DSM-III-R Psychiatric Disorders in the United States: Results From the National Comorbidity Survey. *Archives of General Psychiatry*, 51(1), 8–19. <https://doi.org/10.1001/archpsyc.1994.03950010008002>
- Kuhn, M. (2008). Building predictive models in R using the caret package. *Journal of Statistical Software*, 28(5), 1-26.
- Kuhn, M., & Johnson, K. (2013). *Applied predictive modeling* (Vol. 26). New York: Springer. Chicago
- Liaw, A., & Wiener, M. (2002). Classification and Regression by randomForest. *R News* 2(3), 18-22.
- Matsuki, K., Kuperman, V., & Van Dyke, J. A. (2016). The Random Forests statistical technique: An examination of its value for the study of reading. *Scientific Studies of Reading*, 20(1), 20-33.
- Meehl, P. E. (1990). Why summaries of research on psychological theories are often uninterpretable. *Psychological Reports*, 66, 195-244.

- Mitchell, S. D. (2015). Explaining Complex Behavior. In K. S. Kendler & J. Parnas (Eds.), *Philosophical Issues in Psychiatry: Explanation, Phenomenology, and Nosology* (pp. 19–47). Baltimore: Johns Hopkins University Press
- Morgan, J. N., & Sonquist, J. A. (1963). Problems in the analysis of survey data, and a proposal. *Journal of the American Statistical Association*, 58(302), 415-434.
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251), aac4716.
- Patel, M. J., Khalaf, A., & Aizenstein, H. J. (2016). Studying depression using imaging and machine learning methods. *NeuroImage: Clinical*, 10, 115-123.
- R Core Team (2019). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Ribeiro, J. D., Huang, X., Fox, K. R., Walsh, C. G., & Linthicum, K. P. (2019). Predicting Imminent Suicidal Thoughts and Nonfatal Attempts: The Role of Complexity. *Clinical Psychological Science*, 7, 941-957.
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22(11), 1359-1366.
- Steyerberg, E. W., Harrell Jr, F. E., Borsboom, G. J., Eijkemans, M. J. C., Vergouwe, Y., & Habbema, J. D. F. (2001). Internal validation of predictive models: efficiency of some procedures for logistic regression analysis. *Journal of Clinical Epidemiology*, 54(8), 774-781.

- Steyerberg, E. W., Bleeker, S. E., Moll, H. A., Grobbee, D. E., & Moons, K. G. (2003). Internal and external validation of predictive models: a simulation study of bias and precision in small samples. *Journal of Clinical Epidemiology*, 56(5), 441-447.
- Strobl, C., Malley, J., & Tutz, G. (2009). An introduction to recursive partitioning: rationale, application, and characteristics of classification and regression trees, bagging, and random forests. *Psychological Methods*, 14(4), 323.
- Tantithamthavorn, C., McIntosh, S., Hassan, A. E., & Matsumoto, K. (2016). An empirical comparison of model validation techniques for defect prediction models. *IEEE Transactions on Software Engineering*, 43(1), 1-18.
- van der Ploeg, T., Austin, P. C., & Steyerberg, E. W. (2014). Modern modelling techniques are data hungry: a simulation study for predicting dichotomous endpoints. *BMC Medical Research Methodology*, 14, 1-13.
- Walsh, C. G., Ribeiro, J. D., & Franklin, J. C. (2017). Predicting risk of suicide attempts over time through machine learning. *Clinical Psychological Science*, 5(3), 457-469.
- Yarkoni, T., & Westfall, J. (2017). Choosing prediction over explanation in psychology: Lessons from machine learning. *Perspectives on Psychological Science*, 12(6), 1100-1122.
- Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2), 301-320.
- Zuromski, K. L., Bernecker, S. L., Gutierrez, P. M., Joiner, T. E., King, A. J., Liu, H., ... & Stein, M. B. (2019). Assessment of a risk index for suicide attempts among US Army Soldiers with suicide ideation: analysis of data from the Army Study to Assess Risk and Resilience in Servicemembers (Army STARRS). *JAMA Network Open*, 2(3), e190766-e190766.

Footnotes

1. We use the term *algorithm* to refer to the over-arching statistical algorithm (i.e., random forests or lasso regression) and the term *model* to refer to a trained algorithm with specific tuning parameters (random forests with 500 trees and number of randomly selected predictors of 10).
2. It would be more realistic to include predictors that have extremely small associations with Y (e.g., the crud factor; Meehl, 1990), however, our main goal in the simulation was to study the bias of the validation strategies. For this reason, we wanted to keep the true relationship between X's and Y to be simple and relatively easy to calculate a ground truth performance. Our main concern with the simulation conditions was with respect to both the sample size and number of predictors. See Burke et al., (2019) for a review of machine learning studies in suicide research. Although the review identified a number of studies that used larger sample sizes and number of predictors than we tested in the simulation, we did not believe that higher values for these conditions would significantly change the results, which was confirmed in the simulation results.
3. We do not report the specific values of tuning parameters used for random forests or gradient boosting, as the defaults used vary depending on the simulation condition. For instance, for random forests, *mtry* takes different values depending on the number of predictors. We note that in actual empirical results it is recommended to report the specific values.

Table 1. Areas under the receiver operator characteristic curves (AUC) across models and validation approach in the National Comorbidity Survey (NCS) data.

	GLM	Elastic Net	Gradient Boosting Machine	Random Forest
10-Fold CV	0.77	0.79	0.79	0.79
10 Repeated CV	0.78	0.79	0.79	0.79
Bootstrap (100)	0.77	0.78	0.77	0.78
Optimism Corrected Bootstrap	0.79	0.79	0.93	0.97

Note. Analyses based on work cited in Franklin, 2019 (Huang et al., 2019). GLM = general linear model.

Figure 1

Schematic for k -fold cross-validation. In the diagram, 5-fold CV is demonstrated, where a model is trained on the “Training” portion of the sample, then tested on the “Test” portion, or those observations not used in the training of the model. This is done five times, with each observation being a part of the training sample four times, part of the test sample once.

Figure 2

Elastic net performance while varying the training set size

Figure 3

Elastic net performance while varying the number of noise predictors

Figure 4

Boosting performance while varying the size of the training set

Figure 5

Boosting performance while varying the number of noise predictors

Figure 6

Random forests performance while varying the training set size

Figure 7

Random forests performance while varying the number of noise predictors

Figure 8

Random forests performance with the removal of one “true” predictor across sample size