Quantifying Optimization Bias in Model Evaluation when using Cross-Validation in
Psychological Science: A Monte Carlo Simulation Study

Stage 1 Registered Report

Hannah Moshontz, Gaylen E. Fronk, Sarah J. Sant'Ana, & John J. Curtin
University of Wisconsin Madison

Psychological scientists are increasingly using machine learning to advance applied or basic science goals. Cross-validation is commonly used for training, selecting, and evaluating machine learning algorithms. Although cross-validation is often described as a tool that necessarily prevents overly optimistic estimates of a final model's performance in new data, not all approaches to cross-validation offer accurate evaluation in all data contexts. Prior work in other fields has established that k -fold cross-validation can provide upwardly biased model evaluation. The extent of this optimization bias in data contexts typical of psychological science is not known. This Monte Carlo simulation study characterizes optimization bias in model evaluation when using three different approaches to cross-validation — k -fold cross-validation, k -fold cross-validation with a test set, and nested k -fold cross-validation — in factorial combinations of sample size (100; 500; 1,000), number of predictors (10; 100; 1,000), and outcome distribution (balanced, with a positive case probability of .50; unbalanced, with a positive case probability of .10). For each unique data context, 1,000 simulations were conducted by randomly sampling from multivariate normal predictor distributions and binomial outcome distributions such that predictors were known to have no population-level association with the outcome. We characterized the average, standard deviation, and maximum performance estimates resulting from model evaluation using each approach to cross-validation in each context, identifying biased evaluations as those with an Area Under the Receiver Operating Characteristic Curve (AUC) value greater than .50.

Quantifying Optimization Bias in Model Evaluation when using Cross-Validation in Psychological Science: A Monte Carlo Simulation Study

Psychological scientists are increasingly using machine learning to optimize prediction and advance theory. Machine learning (also called statistical learning) is not a single technique but instead encompasses a range of data analytic tools developed in statistics and computer science and now commonly used in many science and business settings (James et al., 2013; Kuhn & Johnson, 2018). Machine learning typically involves using rigorous model selection and model evaluation approaches to optimize prediction of a meaningful outcome with statistical algorithms (e.g., decision trees, ridge regression, lasso regression, support vector machines). Psychological scientists are increasingly using machine learning (e.g., Takahashi et al., 2020; Wang & Kosinski, 2018; Webb et al., 2020) and recommending that others use machine learning (Kosinski et al., 2016; Orrù et al., 2020; Vieira et al., 2020).

Machine learning is a natural fit for psychological research. In psychology, outcomes of interest reflect many interconnected factors; thoughts, behavior, decisions, and well-being have complex causes. Machine learning methods can accommodate this complexity, allowing researchers to use hundreds of predictors and their interactions and to model non-linear functional forms. Furthermore, machine learning does not require that researchers know which predictors to use or which interactions and functional forms to specify a priori. Instead, researchers can make robust, data-driven decisions, fitting many plausible model configurations and selecting the one that predicts new observations best (i.e., the model that most closely matches the true, population-level data generating process).

Because machine learning models are optimized to predict new observations, these models can be developed to solve challenging prediction problems for potential real-world use.

For example, in clinical psychology, machine learning has been used to predict both suicidal ideation and suicide attempts (Roy et al., 2020; Walsh et al., 2018). In metascience, Altmejd and colleagues (2019) used machine learning to predict whether results of psychology and economics studies would replicate based on each study's sample size, effect size, design, and other characteristics. When accurate, predictive models can be practically useful, especially in clinical settings. For example, models can help match people to the treatments that are expected to work best for them (i.e., precision mental health) or identify risk for adverse events like depressive episodes, so that they can be prevented (Dwyer et al., 2018).

Machine learning is useful not only for its potential applied value but also for developing, refining, and testing theory in psychological science (Bleidorn & Hopwood, 2019; Dwyer et al., 2018; Yarkoni & Westfall, 2017). Machine learning models can identify variables that predict behavior, which can inform theory about causes and moderators of these behaviors. For example, Steele et al. (2018) assessed the incremental predictive contribution of variables derived from fMRI functional connectivity during an inhibitory task after controlling for different combinations of self-reported clinical assessment variables (e.g., age, years of drug use, motivation for change). In doing so, they supported a theoretically meaningful role (i.e., predictive utility above and beyond other clinical variables) for executive inhibitory control in adherence to a substance use disorder treatment program.

Machine learning models can also be used to refine theory or constrain the generalizability of a theory when theory-derived variables fail to predict outcomes they should. For example, Weidman et al. (2020) were not able to predict within-person changes in people's happy mood from acoustic signals in their voices, despite theoretical and empirical consensus that voices communicate and reflect emotion (e.g., through pitch). Failure to predict within-

person changes in happiness from vocal signals was consistent across two samples and a variety of approaches that used, for example, different kinds of vocal recordings (e.g., descriptions of daily events, descriptions of neutral pictures), different outcome variables (e.g., binary and continuous measures), and different statistical algorithms. The authors considered these results along with other empirical work that successfully predicted trait happiness and distinct moods (e.g., happiness vs. other moods) from vocal signals and suggested a refinement of theory: Acoustic qualities of people's voices reflect their mood, but vocal signals reflecting people's day-to-day, temporal changes in single moods are subtle to the point of being undetectable with tools currently available to most psychological scientists.

These and other examples (Al Zoubi et al., 2018; Breen et al., 2019; Delgadillo & Gonzalez Salas Duhne, 2020; Hassanpour et al., 2018; Stephansen et al., 2018) demonstrate that, like any analytic tool at the disposal of psychological scientists, machine learning can support theoretical insights and important practical applications if used appropriately and thoughtfully (Bleidorn & Hopwood, 2019; Dwyer et al., 2018; Yarkoni & Westfall, 2017). Accurate evaluation of model performance is a key component of using machine learning models appropriately and a necessary condition of accomplishing most applied and basic science goals. If a model's estimated performance is overly optimistic (i.e., positively biased), using that model in an applied setting or to support theoretical claims could cause missed intervention opportunities, wasted resources, or other harms.

This Monte Carlo simulation study examines how the approach to cross-validation in machine learning can result in positively biased model evaluation. The next sections provide brief primers on the use of cross-validation for model training, model selection, and model evaluation, first generally and then in the context of three specific, common approaches to cross-

validation. Rather than providing a comprehensive overview, these primers motivate and contextualize the present simulation study. Between and even within fields, different terminology is used to refer to the same machine learning techniques. Thus, for purposes of instruction as well as clarity, we also identify key terms with *italics* and then define or describe them. This terminology was adapted from textbooks by James et al. (2013) and Hastie et al., (2009), who are recognized as long-standing pioneers in machine learning among applied statisticians.

Using Cross-Validation to Develop a Machine Learning Model

Developing a machine learning model typically entails 1) training models with multiple candidate model configurations, 2) assessing each model and selecting a best configuration, and 3) evaluating how well a model with that best configuration will perform with new observations¹. *Candidate model configurations* can vary with respect to the predictors included and/or how predictors are transformed (e.g., with power transformations that affect distributional shape) or combined (e.g., additively, interactively). Model configurations can also vary with respect to the *statistical algorithm* (also called the statistical learning model; e.g., lasso, ridge and elastic net regressions; random forest; neural network; support vector machine) used. The set of candidate model configurations often includes variations of the same statistical algorithm with different *hyperparameter values* that control aspects of the algorithm's operation. For example,

¹ Sometimes there is a fourth stage: training a final model using the best configuration and all available data. This final model is the one that would be used to guide real-world decision-making in applied settings.

the hyperparameter λ controls the degree of regularization in penalized regression algorithms. Crossing variation on predictors, statistical algorithms, and hyperparameter values expands the number of candidate model configurations dramatically; developing a machine learning model can easily involve training thousands of models. In most implementations of machine learning, the number of candidate model configurations nearly ensures that some models will fit the data on which they were trained very well due to chance alone. For this reason, model configurations are assessed and selected on the basis of their relative performance in new data.

Cross-validation (CV) is a class of resampling approaches that enables researchers with one dataset to assess models' performance in new data (i.e., in observations that were not used to train them). These assessments in new data serve important functions in selecting model configurations and evaluating final models. Broadly, CV approaches create multiple sets of data by drawing samples from the full dataset. These different (sub)sets can be used for model training, model selection, and model evaluation.

Training sets are samples used for *model training* (also called model fitting). During model training, models with each candidate model configuration are fit to the training data. For example, during training, model parameters are estimated for regression algorithms, and weights are established for neural network algorithms. Some non-parametric algorithms, like k -nearest neighbors, do not estimate parameters but simply “memorize” the training sets for subsequent predictions.

Validation sets are samples used for *model selection* (or, more accurately, for model configuration selection). During model selection, each (trained) model — one for every candidate model configuration — is used to make predictions for observations in a validation set

that, importantly, does not overlap with the model's training set. On the basis of each model's performance in the validation set, the relatively best model configuration (i.e., the configuration of the model that performs best relative to all other model configurations) is identified and selected.

Test sets are samples used for *model evaluation*. During model evaluation, the performance of the final model is estimated. How well the final model would perform in new data cannot be directly evaluated with CV because the final model is trained on the entire dataset. Further, a final model is sometimes hypothetical because not all researchers create it. So, to estimate the performance of the final model, a model with the selected best configuration is used to make predictions for observations in a test set.

Different approaches to CV vary in how they divide the full dataset into training sets, validation sets, and test sets. Depending on the CV approach, different numbers of training, validation, and test sets are drawn from the full dataset to be used for their respective purposes. Some CV approaches (e.g., k -fold CV) do not have separate test sets but instead use the same sample(s) for selection and evaluation. Critically, both successful theory testing and applied use of machine learning models hinge on the validity of the model evaluation procedure. In the next section, we describe three approaches to CV that are the focus of this simulation study.

Three Approaches to CV

Three common CV approaches that psychologists can implement are 1) k -fold CV, 2) k -fold CV with a test set, and 3) nested k -fold CV. Here, we describe the procedure that each approach uses for model training, selection, and evaluation. Figure 1 depicts key similarities and differences in how these CV approaches divide data for each of these purposes.

K-fold CV

K-fold CV is a basic approach to CV that provides the foundation for other, more complicated approaches to CV. In *k*-fold CV, the full dataset is partitioned into *k* approximately equal-sized, non-overlapping samples referred to as *folds*. The number of folds (*k*) can be as small as two and as large as the sample size of the full dataset (*N*), with typical values being 5 and 10 (James et al., 2013). *K*-fold CV is an iterative process.

To start, the first of the *k* folds is *held out* (i.e., put aside and excluded from model training). The remaining *k*-1 folds serve as the training set: For each candidate model configuration, a model is trained with all observations except those in the first (held-out) fold. This process is iterated *k* times — once for every fold — such that each fold is held out from one training set, and each candidate configuration is used to train *k* models on different, but overlapping, sets of observations.

Then, a single candidate model configuration is selected. First, each fold serves as a validation set for the training set from which it was held out. That is, each trained model is used to make predictions for observations in the fold that was excluded from its training set. The candidate model configuration with the best relative mean performance across the *k* validation sets is selected as the best model configuration. When assessing models (e.g., for the purpose of selecting a best configuration), researchers can quantify performance by using any one of a variety of performance metrics (e.g., accuracy, Area Under the Receiver Operating Characteristic Curve (AUC), mean squared error) that fit their data context and the goals of their research.

Finally, the selected best model configuration is evaluated. *K*-fold CV does not use different validation and test sets. Instead, the performance of the best model configuration (i.e.,

the average performance of models with this configuration across the validation sets) also serves as the estimate of the final model's expected performance in new data. In other words, the same samples serve as validation sets to select the best model configuration and as test sets to evaluate the future performance of the final model using this best configuration in new data. This final model is obtained by re-training a model with the best configuration using the full dataset.

K-fold CV with a Test Set

K-fold CV with a test set is an adaption of k -fold CV that adds an independent test set for model evaluation that is held out from model training and model selection. In the first step of this CV approach, a sample of the full dataset is held out to serve as an independent test set later. This independent test set typically includes 20 - 30% of the total sample size from the full dataset, although recommendations for its size differ depending on data contexts and modeling goals. Once this test set is held out, the same model training and selection processes are used as in k -fold CV using the remaining, held-in data. That is, the remaining data are divided into k folds and used to train models with all candidate model configurations on training sets and select the best model configuration based on the models' performance in validation sets.

Then, for model evaluation, a model with the best configuration is trained on all held-in data (i.e., the 70 – 80% of data remaining after the test set was held out). This retrained model is used to make predictions for observations in the previously unused held-out test set to estimate the expected performance of the final model in new data. The final model is derived by training the best configuration on the full dataset.

Nested CV

Nested CV is an adaptation of k -fold CV with a test set that uses multiple independent test sets for model evaluation. Rather than using just one held-out test set and applying k -fold CV to

the remaining data, nested k -fold CV repeats this procedure multiple times. In nested CV, there are two nested loops for partitioning and holding out folds: an outer loop where held-out folds serve as test sets for model evaluation and inner loops where held-out folds serve as validation sets for model selection just as in k -fold CV.

First, the full dataset is divided into k_{outer} approximately equally-sized folds. Each of these k_{outer} folds serves once as an independent test set to evaluate the model configuration trained and selected with k -fold CV (using k_{inner} folds) applied to the remaining data (i.e., the $k_{\text{outer}} - 1$ outer folds). Nested CV therefore produces k_{outer} best model configurations, which are each evaluated in a different test set, and k_{outer} assessments with which to estimate the performance of the final model in new data. The performance estimate in nested CV is the set of k_{outer} assessments and can be summarized by central tendency and dispersion. The final model is derived by applying k_{inner} -fold CV to the full dataset.

Figure 1. Schematic depiction of model training, selection, and evaluation in three approaches to CV. Blue fill indicates folds used for model training, yellow fill indicates folds used for model selection (or for model selection *and* model evaluation in the case of *k*-fold CV), and green fill indicates folds used only for model evaluation. Dashed lines show the correspondence between test sets and the rest of the data in *k*-fold CV with a test set and nested CV.

Optimization Bias in Model Evaluation of Different CV Approaches

Model evaluation estimates the final model's performance in new data, and the most accurate estimates of model performance in new data come from approaches to CV that use truly new data for model evaluation. K -fold CV with a test set uses an independent test set that has not previously been used in any way for model training or selection. Nested CV uses multiple test sets, while maintaining separation between training sets, validation sets, and test sets. So, although nested CV reuses observations for selection and evaluation, the same observations are never used to select and evaluate any one model.

In contrast, approaches to CV that do not use truly new data for model evaluation can produce *optimization bias*, which is systematic error in estimates of model performance in new data (e.g., Jonathan et al., 2000). K -fold CV reuses observations for selection and evaluation; models are assessed with observations that were previously used to select them. The dual use of the same observations to select and evaluate single models in k -fold CV can produce optimization bias.

Aspects of the data context and the number and type of candidate model configurations affect the magnitude of optimization bias. In some data contexts, this bias can be substantial. K -fold CV can produce performance estimates of models that surpass or far exceed conventional thresholds of acceptable model performance (e.g., the thresholds described by Whelan & Garavan, 2014), when in reality the models will perform no better than chance in new data. Biased model evaluation of k -fold CV has been demonstrated in toy data and simulated data from other research fields (Baumann & Baumann, 2014; Kohavi, 1995; Rao et al., 2008; Stone, 1974). Consequently, research fields with a longer history of using machine learning than psychology (e.g., genetics and chemistry) do not typically use k -fold CV for model evaluation.

Unfortunately, the use of k -fold CV for model evaluation is very common in psychology. Many papers using k -fold CV to train, select, and evaluate machine learning models have been published recently in high-impact psychology journals, including PNAS, Psychological Medicine, Biological Psychiatry, Journal of Personality and Social Psychology, and others (e.g., Ebert et al., 2019; Eichstaedt et al., 2018; Haynos et al., 2020; Kinreich et al., 2019; Moberget et al., 2019; Niemann et al., 2020; Pariyadath et al., 2014; Wang & Kosinski, 2018; Weidman et al., 2020; Youyou et al., 2017). It appears that many researchers, journal editors and reviewers, and likely readers in psychology are either uninformed about the limitations of basic k -fold CV for model evaluation or unaware of better alternative approaches. Apparently, evidence of optimization bias when using k -fold CV has yet to influence psychologists' use of machine learning; perhaps because this evidence has not been published in journals that psychologists read or because the evidence comes from data contexts (e.g., research questions, sample sizes, characteristics of predictors and outcomes) that are too dissimilar from psychological study designs.

Machine learning has the potential to advance both applied and basic science goals in psychology. However, this potential can only be realized if machine learning tools are used appropriately. Accurate model evaluation is necessary to judge if and when a model performs sufficiently well for a proposed application. It is also necessary to test and refine theories. The frequent misuse of k -fold CV for model evaluation will harm psychological science, particularly as machine learning becomes more commonplace. Like other practices that cause models to appear better than they truly are (e.g., undisclosed analytic flexibility; Simmons et al., 2011), biased model evaluation can waste research resources, slow collective progress, and diminish the credibility of psychological science among researchers and the public.

Goals of the Current Simulation Study

The goal of this paper is to quantify optimization bias in model evaluation when using k -fold CV, k -fold CV with a test set, and nested CV in psychological science. To this end, we conducted a Monte Carlo simulation study comparing model performance estimates produced by each approach in data contexts that are more representative of psychological science than extant simulation studies from other fields in terms of sample size and the number of predictors. Compared with many other fields that use machine learning, studies in psychology have relatively smaller sample sizes, which can exacerbate optimization bias (James et al., 2013; Vabalas et al., 2019), but also tend to involve fewer predictors, which can attenuate optimization bias (James et al., 2013). Another factor that affects optimization bias is the balance of outcomes in data contexts with dichotomous outcomes (Forman & Scholz, 2010; López et al., 2014). Dichotomous outcomes are common in psychological science including in clinical studies (e.g. diagnostic status, treatment response or retention), real-time event prediction (e.g. substance use, emotional state), and group membership (e.g. political party, theory-derived sub-groups). In many instances, dichotomous outcomes of interest are unbalanced, and the focus of prediction is the rarer of the two cases (e.g., positive diagnostic status for psychiatric disorders). Outcome balance may interact with other aspects of a data context (N , number of predictors) to affect bias, but the impact of outcome balance has not yet been evaluated in simulation studies in any field.

In order to characterize the potential bias of each CV approach for typical research scenarios that apply machine learning in psychology, our simulation study examined how different combinations of sample size, number of predictors, and outcome distribution affect optimization bias when using each type of CV. We simulated 1000 datasets for each of 18 data

contexts that involve the factorial crossing of the three key dimensions: sample size (N ; ranging from 100 to 1,000), number of predictors (p ; ranging from 10 to 1,000), and outcome distribution (either *balanced*, with equal probability of positive and negative cases, or *unbalanced*, with a .10 probability of positive cases). The results of this study provide relevant, empirical evidence for psychological researchers and readers of research interested in the specific risk of optimization bias presented by common CV approaches in a particular data context.

Our simulation study focuses only on contexts where the predictors are unrelated to the outcome such that final models should not predict outcomes for new observations above chance. This choice reflected the serious concern that optimistically biased model evaluation can lead researchers to believe they have observed systematic relationships where none exist. Optimistic bias in the evaluation of models that perform no better than chance can provide erroneous support for theory-derived hypotheses and/or support the application of models for outcomes they cannot predict. The choice to focus on models with predictors that were unrelated to the outcome also simplified the detection of bias; unbiased model evaluation in this context should indicate chance level prediction for new observations.

In quantifying optimization bias, we used the model performance estimates that each CV approach produced in each data context to characterize the model performance estimates (i.e., the mean across 1,000 simulations, interpreted relative to chance performance), the variance in model performance estimates (i.e., the standard deviation across 1,000 simulations), and the worst-case, maximum bias (i.e., the maximum across 1,000 simulations) of each CV approach in each context. We then assessed whether each approach was biased (i.e., evaluated models as performing better than chance), whether k -fold CV differed from the other approaches in model

evaluation estimates, and whether the magnitude of relative bias in k -fold CV varied as a function of data context.

Disclosures

This study was preregistered [*Links to preregistration*]. Data and scripts are available at [*links to analysis dataset (aggregated), scripts*]. We report how we determined our sample size, all data exclusions, all manipulations, and all measures in the study.

Method

In this simulation study, k -fold CV, k -fold CV with a test set, and nested CV were used to train, select, and evaluate models in 1,000 simulated datasets for each of 18 different data contexts, which varied by sample size, number of predictors and outcome distribution. Below, we describe the data generating procedure, the machine learning procedure, and the simulation study's key outcomes. The simulations and analyses were conducted in R (R Core Team, 2019). The individual simulations were conducted in parallel across many machines by implementing high throughput computing via HTCondor.

Data Generating Procedure

For each unique data context, predictor variables and an outcome variable were sampled using a seed that was unique to that context and simulation, which allowed us to reproduce the same datasets for each CV approach. Predictor variables were uncorrelated with each other and sampled from a multivariate normal distribution with $\mu = 0$, $\sigma = 1$, and $cov = 0$. The dichotomous outcome variable (positive vs. negative) was sampled from a binomial distribution using the `rbinom` function with the probability for the positive case set to either 0.5 or 0.1. This sampling was performed independent of the predictors such that this outcome variable was uncorrelated with the predictors. The exact number of positive cases for the outcome variable in each

simulated dataset varied. To prevent errors later in the modeling process (e.g., instances where no positive cases were available in a training set) for data contexts with $N = 100$ and an unbalanced outcome distribution, we redrew the outcome variable until there were between eight and 12 positive cases. Each data context varied along the following three dimensions.

Sample Size

There were three levels for the sample size (N): 100, 500, and 1,000, a range that captures common sample sizes in psychological science (Bakker et al., 2012). For example, typical sample sizes in social and personality research (Fraley & Vazire, 2014) and clinical research (Reardon et al., 2019) fall between 100 and 500. We included data contexts with $N = 1,000$ because recent machine learning studies in psychological science have used similar sample sizes (Delgadillo & Gonzalez Salas Duhne, 2020; Ebert et al., 2019; Moberget et al., 2019). Sample sizes larger than 1,000 are used in machine learning analyses in psychological science (e.g., $N = 19,412$; Weidman et al., 2020), but we did not include them in this study. Not only are large samples typical in other fields that use machine learning approaches, but also they attenuate optimization bias. Thus, there is a particular need to characterize optimization bias of approaches to CV in small samples (i.e., $N \leq 1,000$; Vabalas et al., 2019; Varoquaux, 2018).

Number of Predictors

There were three levels for the number of predictors (P): 10, 100, and 1,000. The lower limit, $P = 10$, is a small number of predictors for most machine learning applications, but it is more similar to the number of predictors used in psychological science currently. The moderate number of predictors, $P = 100$, is relevant to many potential machine learning applications in psychology, which can incorporate dozens of theoretical variables with a range of operationalizations, covariates, and interactions among predictors. The upper limit, $P = 1,000$, is

more typical of machine learning applications using natural language processing or predictors extracted from neuroimaging signals. Further, although $P = 1,000$ may not be typical of most psychological applications of machine learning, including simulations with this large number of predictors allowed us to assess high-dimensional data contexts in which the number of predictors exceeded the number of participants.

Outcome Distribution

There were two options for the distribution of the dichotomous outcome variable: balanced (i.e., probability of the positive case = 0.50) or unbalanced (i.e., probability of the positive case = 0.10). We choose .10 for the unbalanced context because it approximates prevalence rates of many outcomes in psychological science (e.g., substance use disorders, Hassanpour et al., 2018; high school drop-out rates, Chapman et al., 2011).

Machine Learning Procedure

The R package caret (Kuhn, 2008) was used to develop machine learning models for each of the 18,000 simulated datasets (1,000 simulations for each of the 18 data contexts) using k -fold CV, k -fold CV with a test set, and nested CV.

We set k to 10 for all approaches as $k = 10$ remains one of the most commonly used values (James et al., 2013). For nested CV, both k_{inner} and k_{outer} were set to 10. For k -fold with a test set, we held out 25% of the full dataset as the test set for all data contexts, as the percentage of data held out for model evaluation typically ranges from 10% to 30% (e.g., 10%, Hassanpour et al., 2018; 30%, Reece & Danforth, 2017).

Candidate Model Configurations

The candidate model configurations varied with respect to the predictors that were included, the statistical algorithm that was used, and hyperparameter values for the associated statistical algorithm.

Candidate model configurations varied with respect to the predictors that were included. The predictors (total P based on the data context) were divided into four non-overlapping subsets. Model configurations included one of 15 combinations of these four subsets of predictors (e.g., set 1, set 1 + set 2, set 1 + set 3, etc.). Combining different sets of predictors in this way can be useful in research contexts where predictor sets are derived from different sources and may have different predictive value (e.g., predictors derived from self-report, predictors derived from behavioral responses). Selection of one predictor set over another, or selection of a particular combination of predictor sets, can offer practically useful and theoretically meaningful information.

Candidate model configurations used one of two different statistical algorithms: elastic net regression (implemented in the `glmnet` package; Friedman et al., 2010) and random forests (implemented in the `ranger` package; Wright & Ziegler, 2017). These two algorithms were chosen because they differ with respect to the model type (parametric vs. non-parametric) such that they vary with respect to their flexibility and underlying assumptions regarding the data generating process. In addition, both of these algorithms are well-established, commonly used by psychologists (Haynos et al., 2020; Takahashi et al., 2020), and easy to use in a range of research contexts without a lot of tailoring (i.e., they have good “out of box” performance; Kuhn & Johnson, 2018).

Within each statistical algorithm, candidate model configurations used different hyperparameters. The elastic net algorithm includes two quantitative hyperparameters: λ

and alpha. The ranger implementation of the random forests algorithm includes two hyperparameters that were varied across configurations: `mTry`, which is quantitative; and `splitrule`, which is categorical with two levels (gini and extratrees). A third ranger hyperparameter, `min.node.size`, was set to 1 for all configurations by default by the caret train function. We allowed the caret train function to adaptively select the optimal values for hyperparameters across configurations using its `tuneLength` parameter. We used `tuneLength = 15` for elastic net and `tuneLength = 100` for random forests to provide sufficient variation in hyperparameter values and approximately equal number of combinations of hyperparameter across algorithms.

Study Outcome

We assessed model performance for model evaluation² using the area under the receiver operating characteristic curve (AUC) for each trained model. The AUC is a commonly used performance metric (Kuhn & Johnson, 2018) that is attractive because it is both scale-invariant and classification threshold-invariant (Whelan & Garavan, 2014); it produces estimates of performance that are directly comparable across different scales, and unlike other performance metrics (e.g., model accuracy), AUC produces estimates of performance that are directly comparable whether the outcome is balanced or unbalanced. Thus, model performance metrics in this simulation study are comparable to each other and comparable to models developed by other researchers, regardless of the units used in those studies.

The value of a model's AUC is the probability that the model will predict a higher score for a randomly selected positive case relative to a randomly selected negative case (Fawcett, 2006). Therefore, a model with an AUC of 0.50 cannot discriminate between positive and

² We assessed AUCs for model selection as well. We focus this section on the use of AUCs for model evaluation because that is the study outcome.

negative outcome classes; it predicts these two outcomes at chance. A model with an AUC of 1.0 perfectly discriminates between positive and negative outcome classes; it predicts outcomes without making any errors. By convention, an AUC of 0.60-0.70 indicates weak prediction accuracy, an AUC of 0.70 to 0.85 indicates moderate prediction accuracy, and an AUC more than 0.85 indicates good prediction accuracy (Whelan & Garavan, 2014). In the simulated datasets in this study, predictors were unrelated to the outcome variable, so an unbiased estimate of any model's performance should yield an AUC of 0.50. We quantified optimization bias in this study as any positive deviation in AUC from 0.50, with larger deviations corresponding to more bias.

Analysis Approach

Our analyses were designed first to characterize model performance estimates of each CV approach and then to examine 1) whether each CV approach shows optimization bias, 2) whether optimization bias is greater when using k -fold CV than when using k -fold CV with a test set or nested CV, and 3) whether aspects of data context are associated with differences in optimization bias of k -fold CV (relative to k -fold CV with a test set and nested CV).

We first characterized the model performance estimates, calculating the mean, standard deviation, and maximum AUC across simulations for each approach and in each data context. Then, AUCs for each simulated dataset were analyzed in a general linear model (GLM) with between-subjects regressors for N (100, 500, 1,000), P (10, 100, 1,000), and outcome distribution (balanced, unbalanced) and repeated-measures contrasts for CV approach (k -fold CV, k -fold CV with a test set, nested CV). Helmert codes were used for the between-subjects regressors to compare each level of each factor to the subsequent levels. The exact coding scheme used was

guided by the predicted effect of each factor on optimization bias, so that positive parameter estimates would indicate larger AUCs and more optimization bias. N was coded with two regressors: one that compared simulations with $N = 500$ or $N = 1,000$ to simulations with $N = 100$; and one that compared simulations with $N = 1,000$ to simulations with $N = 500$. P was coded with two regressors: one that compared simulations with $P = 10$ or $P = 100$ to simulations with $P = 1,000$; and one that compared simulations with $P = 10$ to simulations with $P = 100$. Outcome distribution was represented by one regressor that compared simulations with balanced outcomes to those with unbalanced outcomes.

To examine optimization bias for each CV approach, we estimated the model with three different contrasts, one for each approach to CV (i.e., k -fold CV was assigned a value of 1 and the other approaches were assigned a value of 0; k -fold CV with a test set was assigned a value of 1 and the other approaches were assigned a value of 0; nested CV was assigned a value of 1 and the other approaches were assigned a value of 0). For each contrast, the model intercept was the estimate of the mean AUC resulting from the focal approach to CV (i.e., the approach assigned a value of 1). To test for bias, the parameter estimate for the intercept resulting from each contrast was compared against 0.50.

Next, to examine the relative optimization bias of k -fold CV, we estimated the model with a contrast that compared k -fold CV to the average of k -fold CV with a test set and nested CV (i.e., k -fold CV was assigned a value of 1 and k -fold CV with a test set and nested CV were each assigned a value of -0.5). With this coding scheme, the model intercept was the estimated difference between the mean AUC resulting from k -fold CV and the mean AUC resulting from the other two approaches, with positive intercept values indicating a relatively higher mean AUC for k -fold CV (i.e., more optimization bias). To test for a difference in the mean AUC resulting

from k -fold CV vs. the other two CV approaches, the parameter estimate for the intercept resulting from this contrast was compared against 0. Further, within this coding scheme, the parameter estimates represent the effect of each regressor on the relative optimization bias of k -fold CV. For example, the parameter estimate for the outcome distribution regressor estimates the average increase in AUC associated using k -fold CV for evaluation (vs. other CV approaches) in simulations with unbalanced versus balanced outcomes. To test for the moderating effects of data context on relative optimization bias of k -fold CV, the parameter estimates for each Helmert-coded regressor were compared against 0. Analyses were preregistered at [*osf links to the pre-registration and to pre-registered analysis code*]. Additional exploratory analyses examined differences in bias when using k -fold CV with a test set compared to nested CV, and moderation of this difference by data context.

Sample Size and Power

The sample size for this study was determined in advance and was a consequence of the choice to conduct 1,000 simulations for each of the 18 data contexts. This number of simulations struck a balance between feasibility in computation time and power. The resulting sample size (i.e., the 18,000 simulated datasets evaluated by each of three CV approaches) was expected to provide sufficient power to detect small but practically important differences in bias among the CV approaches. A priori sensitivity power analyses (conducted using the `lmSupport` package (Curtin, 2015) in R, which calculates power based on the approach outlined by Cohen (1988)) confirmed our power to detect small differences between k -fold CV vs. the other two CV approaches within our proposed GLM. These analyses determined that we had $> 99.9\%$ power to detect a small effect size (Cohen's $d = 0.2$) for the contrast of k -fold CV vs. the other two CV approaches in the previously specified GLM with 18,000 simulated datasets, repeated measures

on CV approach, and a two-tailed $\alpha = 0.05$. Furthermore, the minimum effect size that we can detect for this contrast with $\sim 90\%$ power using these same parameters is Cohen's $d = 0.05$.

Results

Model performance estimates of each CV approach are presented in Figure 2. *[Figure 2 will show a violin plot overlaid on a box plot depicting AUC by CV approach (k-fold CV vs. the average of the other two approaches.) Mean, standard deviation, and maximum values will be reported in the caption.]* The mean, standard deviation, and maximum AUC values for each approach to CV in each data context are reported in Supplemental Table 1.

As shown in Table 1 there was *[sufficient/insufficient]* evidence that model evaluation AUCs were different from 0.50, suggesting *[the presence/absence]* of optimization bias when using *[approaches to CV with significant t-test values comparing intercepts to 0.50; repeat as needed to accommodate both significant and non-significant results]* *[Table 1 will show the value, CI, and SE of the parameter estimates for the intercept, and, df, t, and p-values, for the tests all three intercept parameter estimates against 0.5.]*

As shown in Table 2, *k*-fold CV *[differed/did not differ]* from *k*-fold CV with a test set and nested CV in model evaluation AUCs, suggesting that *[interpret the parameter estimate, indicate whether k-fold CV is relatively more or less biased than the other approaches]*. *[Table 2 will show all parameter estimates from the GLM for the intercept and regressors, along with each of their CI, SE, df, t, and p-values]* As shown in Table 2, *[characterize the general pattern of results related to data context moderation.]* As shown in Figure 3, for *N*, optimization bias with *k*-fold CV *[characterize general trend (e.g., “was greatest in data contexts with smaller N.”)]* *[Interpret the specific parameter estimates for Helmert-coded regressors for N.]* *[Figure 3*

will show a violin plot overlaid on a box plot depicting AUC by N (with three levels) and CV approach (with two levels: k-fold CV and the average of the other two approaches.)] As shown in Figure 4, for P, optimization bias with k-fold CV [repeat the same structure of reporting results for P as was done for N.] [Figure 4 will show a violin plot overlaid on a box plot depicting AUC by P (with 3 levels) and CV approach (with two levels: k-fold CV and the average of the other two approaches.)] As shown in Figure 5, for outcome distribution, optimization bias with k-fold CV [differed/ did not differ] as a function of balanced versus unbalanced outcomes. [Interpret the parameter estimate for outcome balance.] [Figure 5 will show a violin plot overlaid on a box plot depicting AUC by outcome distribution and CV approach (with two levels: k-fold CV and the average of the other two approaches.)]

[Summary of labeled exploratory analyses, focused on differences between k-fold CV with a test set and nested CV will be provided here.]

Author Contributions: J.J. Curtin generated the idea for this study. H. Moshontz, J.J. Curtin, G.E. Fronk, and S.J. Sant’Ana jointly designed the study. H. Moshontz and J.J. Curtin wrote the manuscript. G.E. Fronk, and S.J. Sant’Ana critically edited it. All authors approved the final submitted version of the manuscript.

Conflicts of Interest: The author(s) declare that there were no conflicts of interest with respect to the authorship or the publication of this article.

Funding: This research was supported by two grants to Curtin from the NIH (R01 DA047315; R01 AA024391). This research was performed using the compute resources and assistance of the UW-Madison Center For High Throughput Computing (CHTC) in the Department of Computer Sciences. The CHTC is supported by UW-Madison, the Advanced Computing Initiative, the Wisconsin Alumni Research Foundation, the Wisconsin Institutes for Discovery, and the National Science Foundation, and is an active member of the Open Science Grid, which is supported by the National Science Foundation and the U.S. Department of Energy's Office of Science.

Supplemental Material: [*Link to supplemental table.*]

Prior versions: The submitted version of this manuscript is archived at [*link to PsyArXiv preprint.*]

References

- Al Zoubi, O., Ki Wong, C., Kuplicki, R. T., Yeh, H.-W., Mayeli, A., Refai, H., Paulus, M., & Bodurka, J. (2018). Predicting Age From Brain EEG Signals-A Machine Learning Approach. *Frontiers in Aging Neuroscience*, 10, 184.
<https://doi.org/10.3389/fnagi.2018.00184>
- Altmejd, A., Dreber, A., Forsell, E., Huber, J., Imai, T., Johannesson, M., Kirchler, M., Nave, G., & Camerer, C. (2019). Predicting the replicability of social science lab experiments. *PLOS ONE*, 14(12), e0225826. <https://doi.org/10.1371/journal.pone.0225826>
- Bakker, M., van Dijk, A., & Wicherts, J. M. (2012). The Rules of the Game Called Psychological Science. *Perspectives on Psychological Science*, 7(6), 543–554.
<https://doi.org/10.1177/1745691612459060>
- Baumann, D., & Baumann, K. (2014). Reliable estimation of prediction errors for QSAR models under model uncertainty using double cross-validation. *Journal of Cheminformatics*, 6(1).
<https://doi.org/10.1186/s13321-014-0047-1>
- Bleidorn, W., & Hopwood, C. J. (2019). Using Machine Learning to Advance Personality Assessment and Theory. *Personality and Social Psychology Review*, 23(2), 190–203.
<https://doi.org/10.1177/1088868318772990>
- Breen, M. S., Thomas, K. G. F., Baldwin, D. S., & Lipinska, G. (2019). Modelling PTSD diagnosis using sleep, memory, and adrenergic metabolites: An exploratory machine-learning study. *Human Psychopharmacology: Clinical and Experimental*, 34(2), e2691.
<https://doi.org/10.1002/hup.2691>

- Chapman, C., Laird, J., Ifill, N., & Kewal Ramani, A. (2011). *Trends in High School Dropout and Completion Rates in the United States: 1972-2009*. [Compendium Report. NCES 2012-006.]. National Center for Education Statistics.
- Cohen, J. (1988). *Statistical Power Analysis for the Behavioral Sciences* (2 edition). Routledge.
- Curtin, J. J. (2015). *lmSupport: Support for Linear Models* (2.9.2) [Computer software].
<https://cran.r-project.org/web/packages/lmSupport/index.html>
- Delgadillo, J., & Gonzalez Salas Duhne, P. (2020). Targeted prescription of cognitive–behavioral therapy versus person-centered counseling for depression using a machine learning approach. *Journal of Consulting and Clinical Psychology*, 88(1), 14–24.
<https://doi.org/10.1037/ccp0000476>
- Dwyer, D. B., Falkai, P., & Koutsouleris, N. (2018). Machine Learning Approaches for Clinical Psychology and Psychiatry. *Annual Review of Clinical Psychology*, 14(1), 91–118.
<https://doi.org/10.1146/annurev-clinpsy-032816-045037>
- Ebert, D. D., Buntrock, C., Mortier, P., Auerbach, R., Weisel, K. K., Kessler, R. C., Cuijpers, P., Green, J. G., Kiekens, G., Nock, M. K., Demyttenaere, K., & Bruffaerts, R. (2019). Prediction of major depressive disorder onset in college students. *Depression and Anxiety*, 36(4), 294–304. <https://doi.org/10.1002/da.22867>
- Eichstaedt, J. C., Smith, R. J., Merchant, R. M., Ungar, L. H., Crutchley, P., Preoțiu-Pietro, D., Asch, D. A., & Schwartz, H. A. (2018). Facebook language predicts depression in medical records. *Proceedings of the National Academy of Sciences*, 115(44), 201802331. <https://doi.org/10.1073/pnas.1802331115>
- Fawcett, T. (2006). An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8), 861–874. <https://doi.org/10.1016/j.patrec.2005.10.010>

- Forman, G., & Scholz, M. (2010). Apples-to-apples in cross-validation studies: Pitfalls in classifier performance measurement. *ACM SIGKDD Explorations Newsletter*, 12(1), 49–57. <https://doi.org/10.1145/1882471.1882479>
- Fraley, R. C., & Vazire, S. (2014). The N-Pact Factor: Evaluating the Quality of Empirical Journals with Respect to Sample Size and Statistical Power. *PLoS ONE*, 9(10). <https://doi.org/10.1371/journal.pone.0109019>
- Friedman, J. H., Hastie, T., & Tibshirani, R. (2010). Regularization Paths for Generalized Linear Models via Coordinate Descent. *Journal of Statistical Software*, 33(1), 1–22. <https://doi.org/10.18637/jss.v033.i01>
- Hassanpour, S., Tomita, N., DeLise, T., Crosier, B., & Marsch, L. A. (2018). Identifying substance use risk based on deep neural networks and Instagram social media data. *Neuropsychopharmacology*, 1. <https://doi.org/10.1038/s41386-018-0247-x>
- Hastie, T., Tibshirani, R., & Friedman, J. H. (2009). *The elements of statistical learning: Data mining, inference, and prediction* (2nd ed). Springer.
- Haynos, A. F., Wang, S. B., Lipson, S., Peterson, C. B., Mitchell, J. E., Halmi, K. A., Agras, W. S., & Crow, S. J. (2020). Machine learning enhances prediction of illness course: A longitudinal study in eating disorders. *Psychological Medicine*, 1–11. <https://doi.org/10.1017/S0033291720000227>
- James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An Introduction to Statistical Learning: With Applications in R* (7th ed.). Springer-Verlag. <https://doi.org/10.1007/978-1-4614-7138-7>

- Jonathan, P., Krzanowski, W. J., & McCarthy, W. V. (2000). On the use of cross-validation to assess performance in multivariate prediction. *Statistics and Computing*, 10(3), 209–229. <https://doi.org/10.1023/A:1008987426876>
- Kinreich, S., Meyers, J. L., Maron-Katz, A., Kamarajan, C., Pandey, A. K., Chorlian, D. B., Zhang, J., Pandey, G., Subbie-Saenz de Viteri, S., Pitti, D., Anokhin, A. P., Bauer, L., Hesselbrock, V., Schuckit, M. A., Edenberg, H. J., & Porjesz, B. (2019). Predicting risk for Alcohol Use Disorder using longitudinal data with multimodal biomarkers and family history: A machine learning study. *Molecular Psychiatry*, 1–9. <https://doi.org/10.1038/s41380-019-0534-x>
- Kohavi, R. (1995). A study of cross-validation and bootstrap for accuracy estimation and model selection. *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, 2, 1137–1143.
- Kosinski, M., Wang, Y., Lakkaraju, H., & Leskovec, J. (2016). Mining big data to extract patterns and predict real-life outcomes. *Psychological Methods*, 21(4), 493. <https://doi.org/10.1037/met0000105>
- Kuhn, M. (2008). Building Predictive Models in R Using the caret Package. *Journal of Statistical Software*, 28(1), 1–26. <https://doi.org/10.18637/jss.v028.i05>
- Kuhn, M., & Johnson, K. (2018). *Applied Predictive Modeling* (1st ed. 2013, Corr. 2nd printing 2018 edition). Springer.
- López, V., Fernández, A., & Herrera, F. (2014). On the importance of the validation technique for classification with imbalanced datasets: Addressing covariate shift when data is skewed. *Information Sciences*, 257, 1–13. <https://doi.org/10.1016/j.ins.2013.09.038>

- Moberget, T., Alnæs, D., Kaufmann, T., Doan, N. T., Córdova-Palomera, A., Norbom, L. B., Rokicki, J., van der Meer, D., Andreassen, O. A., & Westlye, L. T. (2019). Cerebellar Gray Matter Volume Is Associated With Cognitive Function and Psychopathology in Adolescence. *Biological Psychiatry*, 86(1), 65–75.
<https://doi.org/10.1016/j.biopsych.2019.01.019>
- Niemann, U., Brueggemann, P., Boecking, B., Mazurek, B., & Spiliopoulou, M. (2020). Development and internal validation of a depression severity prediction model for tinnitus patients based on questionnaire responses and socio-demographics. *Scientific Reports*, 10(1), 4664. <https://doi.org/10.1038/s41598-020-61593-z>
- Orrù, G., Monaro, M., Conversano, C., Gemignani, A., & Sartori, G. (2020). Machine Learning in Psychometrics and Psychological Research. *Frontiers in Psychology*, 10.
<https://doi.org/10.3389/fpsyg.2019.02970>
- Pariyadath, V., Stein, E. A., & Ross, T. J. (2014). Machine learning classification of resting state functional connectivity predicts smoking status. *Frontiers in Human Neuroscience*, 8.
<https://doi.org/10.3389/fnhum.2014.00425>
- R Core Team. (2019). *R: A language and environment for statistical computing*. <https://www.R-project.org/>
- Rao, R. B., Fung, G., & Rosales, R. (2008). On the Dangers of Cross-Validation. An Experimental Evaluation. In C. Apte, H. Park, K. Wang, & M. J. Zaki (Eds.), *Proceedings of the 2008 SIAM International Conference on Data Mining* (Vol. 1–0, pp. 588–596). Society for Industrial and Applied Mathematics.
<https://doi.org/10.1137/1.9781611972788.54>

- Reardon, K. W., Smack, A. J., Herzhoff, K., & Tackett, J. L. (2019). An N-pact factor for clinical psychological research. *Journal of Abnormal Psychology*, 128(6), 493–499. <https://doi.org/10.1037/abn0000435>
- Reece, A. G., & Danforth, C. M. (2017). Instagram photos reveal predictive markers of depression. *EPJ Data Science*, 6(1), 1–12. <https://doi.org/10.1140/epjds/s13688-017-0110-z>
- Roy, A., Nikolitch, K., McGinn, R., Jinah, S., Klement, W., & Kaminsky, Z. A. (2020). A machine learning approach predicts future risk to suicidal ideation from social media data. *Npj Digital Medicine*, 3(1), 1–12. <https://doi.org/10.1038/s41746-020-0287-6>
- 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. <https://doi.org/10.1177/0956797611417632>
- Steele, V. R., Maurer, J. M., Arbabshirani, M. R., Claus, E. D., Fink, B. C., Rao, V., Calhoun, V. D., & Kiehl, K. A. (2018). Machine Learning of Functional Magnetic Resonance Imaging Network Connectivity Predicts Substance Abuse Treatment Completion. *Biological Psychiatry. Cognitive Neuroscience and Neuroimaging*, 3(2), 141–149. <https://doi.org/10.1016/j.bpsc.2017.07.003>
- Stephansen, J. B., Olesen, A. N., Olsen, M., Ambati, A., Leary, E. B., Moore, H. E., Carrillo, O., Lin, L., Han, F., Yan, H., Sun, Y. L., Dauvilliers, Y., Scholz, S., Barateau, L., Hogl, B., Stefani, A., Hong, S. C., Kim, T. W., Pizza, F., ... Mignot, E. (2018). Neural network analysis of sleep stages enables efficient diagnosis of narcolepsy. *Nature Communications*, 9(1), 5229. <https://doi.org/10.1038/s41467-018-07229-3>

- Stone, M. (1974). Cross-Validatory Choice and Assessment of Statistical Predictions. *Journal of the Royal Statistical Society. Series B (Methodological)*, 36(2), 111–147.
- Takahashi, Y., Ueki, M., Yamada, M., Tamiya, G., Motoike, I. N., Saigusa, D., Sakurai, M., Nagami, F., Ogishima, S., Koshiba, S., Kinoshita, K., Yamamoto, M., & Tomita, H. (2020). Improved metabolomic data-based prediction of depressive symptoms using nonlinear machine learning with feature selection. *Translational Psychiatry*, 10(1), 1–12. <https://doi.org/10.1038/s41398-020-0831-9>
- Vabalas, A., Gowen, E., Poliakoff, E., & Casson, A. J. (2019). Machine learning algorithm validation with a limited sample size. *PLOS ONE*, 14(11), e0224365. <https://doi.org/10.1371/journal.pone.0224365>
- Varoquaux, G. (2018). Cross-validation failure: Small sample sizes lead to large error bars. *NeuroImage*, 180, 68–77. <https://doi.org/10.1016/j.neuroimage.2017.06.061>
- Vieira, S., Garcia-Dias, R., & Lopez Pinaya, W. H. (2020). Chapter 19—A step-by-step tutorial on how to build a machine learning model. In A. Mechelli & S. Vieira (Eds.), *Machine Learning* (pp. 343–370). Academic Press. <https://doi.org/10.1016/B978-0-12-815739-8.00019-5>
- Walsh, C. G., Ribeiro, J. D., & Franklin, J. C. (2018). Predicting suicide attempts in adolescents with longitudinal clinical data and machine learning. *Journal of Child Psychology and Psychiatry*, 59(12), 1261–1270. <https://doi.org/10.1111/jcpp.12916>
- Wang, Y., & Kosinski, M. (2018). Deep neural networks are more accurate than humans at detecting sexual orientation from facial images. *Journal of Personality and Social Psychology*, 114(2), 246–257. <https://doi.org/10.1037/pspa0000098>

- Webb, C. A., Cohen, Z. D., Beard, C., Forgeard, M., Peckham, A. D., & Björgvinsson, T. (2020). Personalized prognostic prediction of treatment outcome for depressed patients in a naturalistic psychiatric hospital setting: A comparison of machine learning approaches. *Journal of Consulting and Clinical Psychology*, 88(1), 25–38. <https://doi.org/10.1037/ccp0000451>
- Weidman, A. C., Sun, J., Vazire, S., Quoidbach, J., Ungar, L. H., & Dunn, E. W. (2020). (Not) hearing happiness: Predicting fluctuations in happy mood from acoustic cues using machine learning. *Emotion*, 20(4), 642–658. <https://doi.org/10.1037/emo0000571>
- Whelan, R., & Garavan, H. (2014). When Optimism Hurts: Inflated Predictions in Psychiatric Neuroimaging. *Biological Psychiatry*, 75(9), 746–748. <https://doi.org/10.1016/j.biopsych.2013.05.014>
- Wright, M. N., & Ziegler, A. (2017). ranger: A Fast Implementation of Random Forests for High Dimensional Data in C++ and R. *Journal of Statistical Software*, 77(1). <https://doi.org/10.18637/jss.v077.i01>
- Yarkoni, T., & Westfall, J. (2017). Choosing Prediction Over Explanation in Psychology: Lessons From Machine Learning. *Perspectives on Psychological Science*, 12(6), 1100–1122. <https://doi.org/10.1177/1745691617693393>
- Youyou, W., Stillwell, D., Schwartz, H. A., & Kosinski, M. (2017). Birds of a Feather Do Flock Together: Behavior-Based Personality-Assessment Method Reveals Personality Similarity Among Couples and Friends. *Psychological Science*, 28(3), 276–284. <https://doi.org/10.1177/0956797616678187>