Practical Foundations of Machine Learning for Addiction Research.

Part II. Workflow and use cases.

Pablo Cresta Morgado ^{1,*} , Martín Carusso ¹, Laura Alonso Alemany ² and Laura Acion ^{1,3}

¹ Instituto de Cálculo, FCEyN, Universidad de Buenos Aires - CONICET, Argentina

² FaMAF, Universidad Nacional de Córdoba, Argentina

³ Department of Psychiatry, University of Iowa, Iowa, USA

* Correspondence: pablo_crestam@hotmail.com; Address: Intendente Güiraldes 2160, Ciudad Universitaria. (C1428EGA) Buenos Aires, Argentina.

Abstract: In a continuum with applied statistics, machine learning offers a wide variety of tools to explore, analyze, and understand addiction data. These tools include algorithms that can leverage useful information from data to build models. These models are capable of addressing different scientific problems. In this second part of this two-part machine learning review, we develop how to apply machine learning methods. We explain the main limitations of machine learning approaches and ways to address them. Like other analytical tools, machine learning methods require careful implementation to carry out a reproducible and transparent research process with reliable results. This review describes a helpful workflow to guide the application of machine learning. This workflow has several steps: study design, data collection, data pre-processing, modeling, and communication. How to train, validate and test a model, detect and characterize overfitting, and determine an adequate sample size are some of the key issues to handle when applying machine learning. We also illustrate the process and particular nuances with examples of how researchers in addiction have applied machine learning techniques with different goals, study designs, or data sources.

Keywords: machine learning, data science, artificial intelligence, addiction research methods, data analysis

1. Introduction

Machine learning is a powerful toolbox that assembles a wide diversity of methods and techniques for data analysis (1). It can improve results in several science fields, including biomedicine (2). Addiction research has applied machine learning to address various scientific questions, such as finding individuals at risk of an overdose, identifying participants with a substance use disorder (SUD), or connecting language production to SUD (3–5). Moreover, the flexibility of machine learning tools allows us to apply them to different types of data (e.g., structured data from surveys, images, or tweets), coming from a broad set of data sources (e.g., electronic health records, surveys, neuroimaging, or social media) and different study designs (6–8). In every instance, machine learning can offer helpful tools to improve how data are understood and analyzed. Machine learning can be seen as a natural extension of traditional statistics (9) that can learn from data to improve their performance resolving a task (10). There are a wide variety of these methods, usually classified as supervised (e.g., linear or logistic regression, random forests, and artificial neural networks) and unsupervised (e.g., clustering techniques or principal component analysis) (1,11). There is a broad amount of evidence about machine learning capacities, methods, and techniques in the addiction research field (1). Like with any other extensive set of analytical tools, machine learning requires careful implementation. Researchers must address some issues to avoid obtaining poor results, such as data quality, sample size, overfitting, and how we train and evaluate models. Machine learning implementation needs an adequate scientific framework, explicit assumptions, and an auditable and reproducible workflow. This article is the second part of a two-part review of machine learning in addiction research. In the first part, we introduced foundational concepts of

Citation: Cresta Morgado, P.; Carusso, M.; Alonso Alemany, L.; Acion, L. Machine Learning for Addiction Research. *Preprints* 2021, 1, 0. <https://doi.org/>

machine learning, followed by a description of tools and available methods. This second part examines critical considerations for applying and integrating these methods in a reliable workflow, including addiction research examples. Our supplementary material goes deeper into addiction research published studies that applied different parts of our proposed workflow. Although each review part stands alone, this second part assumes an understanding of some of the first part's definitions. Hence, if the reader is not familiar with machine learning terminology, we recommend starting the reading at Part I. We provide a glossary that defines the key concepts (in italic in the text) used throughout the review as supplementary material for further guidance.

2. Some general considerations when faced with a rich methodological menu

Machine learning provides a broad palette of methods and techniques to address a problem, which may be overwhelming. Hence we start by outlining some of the most crucial decisions regarding the machine learning process and elaborate on them in the rest of the review. Even if time or resources do not allow exploring all the following questions thoroughly, it is beneficial to keep a close check on well-known caveats of methods and data configurations, like bias toward a majority class or the tendency to overfit. A sound methodology involves noting them at the moment of design and monitoring them at various points during the experimental process, as these aspects are possibly the cause of unexpected results. First and foremost, at the research design step, no matter how innovative the methods will be used, researchers must consider if the research should be conducted, its direct or indirect ethical implications, and if the research is worth the investment of any resources.

2.1. Data-related considerations

While the following considerations are well-known to researchers used to working with data, it is easy to lose track of them in the newly added complexity of machine learning. During the research design stage, researchers need to decide which data to collect for their description or prediction goals and ways to obtain data that adequately represent the theoretical framework that generated their research questions. Adequate sample size planning is also crucial, even for prediction tasks. Researchers should consider possible biases and excessive variability during data collection and ways to avoid or mitigate them. Strategies for mitigation may range from acquiring complementary data that allows recovering the actual phenomenon, correcting the error, applying sub-sampling or oversampling in the datasets' constitution, applying feature modifications that equalize feature importance, or applying modifications to models to correct tendencies that mischaracterize the phenomenon under study. All of these strategies are a source of potential errors that require close supervision and documentation.

2.2. Modeling considerations

First of all, researchers must consider which classification or regression goals are feasible, given the data at hand. When the data is insufficient to obtain a reliable model for the given target, gathering more data or complementary features may be needed. Some cases may require changing the target to a more reliable proxy research question related to the original research question. After choosing an algorithm suitable for the task, it is essential to ensure models do not obscure relevant data properties. For example, the model may be misrepresenting minority classes, guided primarily by more prominent features or outliers and not by features or values that may be more relevant to the problem at hand. Also, overgeneralization and overfitting are very pervasive flaws of machine learning solutions. Many methodological mistakes happen during the model evaluation process, like information leakage from training data. Using over-simplistic metrics may reflect only one aspect of the outcome (like overall accuracy) but fail to show errors for different subgroups or results' variance. Complementary metrics are needed to reflect relevant aspects of the performance of the model. If flaws become evident, steps to address them are essential.

One may tweak the model's parameters, rethink feature pre-processing, or even change the data acquisition process. Measuring the impact of any of these modifications in the results is critical, modifying evaluation metrics for more insightful assessment, if necessary. Thus, showing whether and how differences between approaches are significantly different is crucial and is especially relevant when comparing complex, costly methods to baseline or more straightforward approaches. In a nutshell, researchers often face complex research objectives that require many, often arbitrary, decisions because most data analysis aspects do not have clear answers. In what follows, we describe a pragmatic and iterative machine learning workflow in detail, providing ways to address some of the considerations raised thus far.

3. Typical analytical workflow including machine learning

A typical workflow for applying machine learning techniques to a problem consists of the same stages as any other data analysis process: designing the study, collecting and pre-processing or curating the data, modeling the data, evaluating the model, and communicating the results as depicted in Figure 1. In part I of this review (1), we introduced machine learning methods and techniques and described their strengths and limitations. We will now focus on the intricacies of bringing these methods into practice, highlighting known issues

3.1. Design

The driving aim of any study is its scientific question. Next, the analytic goal (e.g., prediction, dimension reduction) needs to be defined to evaluate how that goal is achieved. These aspects require all the knowledge available about the problem that may be of use for the targeted outcome, like the covariates and confounders involved (12). Before data collection, it is essential to include all of these aspects in an analysis plan. Relevant variables are defined at this stage, but it is also essential to plan the variable selection strategy that will be implemented afterward (see pre-processing section). Sample size calculation is critical for any study that aims to test hypotheses. Reliable prediction models also require sample size calculations, more so if we take into account that machine learning models can have numerous parameters to be learned from the data. Small sample sizes can result in parameter estimates that will vary a lot for different samples. To address this, Riley and colleagues (13) proposed four steps (and an accompanying R package) to calculate the sample size required for clinical predictive models (13). In many cases, it may become apparent that the intended research question cannot be reliably answered given the information in the available data and the limitations of the machine learning approach. This issue may come up at the design phase, but it is often clearer at the analysis stage. If that is the case, more or different data may be gathered to obtain more information and reliably address the research question. In other cases, one may opt to redefine the research question, posing a proxy research question related to the original research question but can be addressed using the available data. For example, a research question may be to determine whether a person has an addiction to a given substance, but the proxy question that is finally addressed by the machine learning experiment is to determine whether that person will undergo a binge episode in the following month. There are several high-quality guidelines concerning the must-haves of most common health research designs mainly developed by the EQUATOR Network (14). It is a common misconception to consider EQUATOR Network guidelines only at the study report stage. Using them from the study design stage avoids omitting study's key components. These guidelines are a consensus that results from multidisciplinary work and provides minimum requirements that make research results reproducible and helpful to advance knowledge. There are several of these guidelines developed taking machine learning methods into account. In addition to the TRIPOD statement (15,16) for prediction models, Luo et al. propose a guideline that focuses on using machine learning to build predictive models in biomedical research and to minimize often-omitted items when applying machine learning (17).

Figure 1. Schematic workflow for applying machine learning techniques.

More recently, the CONSORT statement was extended for clinical trials using artificial intelligence interventions (18). Two other guidelines, focused on machine learning and artificial intelligence (i.e., TRIPOD-ML and STARD-AI), are currently under development (19,20). Although the original version of these guidelines (TRIPOD statement (15) and STARD (21)) are useful for machine learning studies, these new initiatives aim to establish methodological standards tailored to machine learning uses by harmonizing terminology and addressing common issues. Differing terminology between applied statistics and machine learning is a well-known issue and the rationale for providing a Glossary for this review (see supplementary material). Machine learning algorithms learn from data. Therefore low-quality data will result in poor models. Data quality is strongly related to its collection, being this a critical step because analytical techniques fall short of reconstructing

missing information or detecting errors in the data. Aiming at high-quality data collection starts at the study design stage.

3.2. Pre-processing

During data collection, the data curation stage involves exploring data to detect problems early. Once data collection is complete, while data exploration continues, comes the need for data cleaning and, in some cases, transformation to facilitate the generation of a suitable model for the research question. As is shown in Figure 1, the data curation stage is iterative and closely related to the definition of the scientific question and the more concrete machine learning task. Detail-oriented, careful, and well-documented pre-processing is necessary because it profoundly affects final results (12). Exploratory data analysis is a great starting point and very informative. Explorative plots can show data quality issues and outliers. Summary statistics (e.g., mean, median, standard deviation) help double-check the data and even find missing values. Under the umbrella of exploratory analyses, unsupervised techniques can reveal data patterns detecting observation groups where none were expected or unveil new problems. They are often helpful in detecting the importance of existing features or creating new features, such as the cluster to which an observation belongs or the value of a data point for the first component in the principal component analysis. Once the data are clean, it is time to deal with missing data. Both theory and empirical evidence show decreased performance if missing data are not completed (22). There are several methods to generate plausible synthetic values, commonly referred to as imputations. The simplest methods fill missing instances with the variable's mean or median. Other, more accurate and complex techniques involve using supervised machine learning algorithms to fill missing values. These methods set the incomplete column as a dependent variable, use the remaining variables to predict the missing values, and simultaneously handle bias sources (i.e., incompleteness, measurement error, and confounders) (22). For instance, Jing et al. (7), in a study to identify youths at risk of having a SUD, applied the k-nearest neighbors algorithm to impute missing data. The authors' rationale was that values extrapolated from the most similar cases could be the best substitute for missing information. Curation also involves feature engineering or variable transformation and variable selection. This step aims to transform the original variables into features that allow modeling the problem better (23), comprising scaling, feature interactions, categorical variable coding (i.e., creation of dummy variables or one-hot encoding), and dimensionality reduction (e.g., using principal component analysis (1)). Variable selection methods aim to keep all relevant predictors for a problem while minimizing bias and variance. Variable selection helps to minimize degrees of freedom for inference tasks, lower prediction variability, and estimate unbiased effects. Variable selection should always involve expert knowledge in the field before recurring to automated methods such as lasso, random forests, or combinations of supervised machine learning algorithms and metrics to assist the variable selection step. Any form of variable selection, creation or transformation, including, for example, linear transformations, can be formally viewed as an embedding, a projection from an initial feature space to a final feature space. The objective of this projection is that observations become more clearly distinguishable in the final model space. Specific machine learning techniques, like neural embeddings, can also infer such projections. This area of machine learning, known as representation learning (1). Curation is often considered more an art than a science due to the many arbitrary decisions required (24,25). That is why detailed, rigorous documentation and transparency about this stage is paramount to the validity and reproducibility of any analysis, even when it does not include most of the machine learning techniques mentioned in previous sections. To this end, The Turing Way guides explain and provide tools and procedures that facilitate these tasks for addiction researchers (26).

3.3. Modeling

There is no simple rule to select the correct method to address a question. Occam's razor recommends starting from the simplest approaches. Thus, best practices recommend starting with the lowest complexity models (e.g., linear models) and use them to establish a baseline benchmark. When more complexity is needed, researchers may scale up to more sophisticated methods, often at the cost of interpretability and with the risk of overfitting training cases. Flexible methods tend to be sensitive to minor variations in the data; thus, they usually require a larger number of cases to identify patterns correctly. Finding the best approach often involves fitting several different models and model parameterizations, as is shown in the examples presented in the supplementary material. Exploring different model families (e.g., logistic regression, random forests, Bayesian methods) and then parameterizing the most promising one is a common way to pursue at this stage. Good practices at this stage include having a rationale for exploring each model family.

3.3.1. Model Performance Evaluation

As shown in Figure 1, modeling is iterative and may include discovering problems in the dataset that were not detected during the pre-processing stage or requiring additional feature engineering to improve model fitting. The final goal of a model is to achieve a good bias-variance trade-off. For predictive models, a good tradeoff results in a good performance with unseen data points. When model parameter estimation is of interest, a good trade-off results in unbiased parameter estimates with low variability, which ultimately results in having good predictive performance with unseen data. Thus, model performance requires using a different dataset from the one used for building the model. For evaluation purposes, we divide annotated datasets into disjoint subsets to fulfill three different roles: the training set, from which the model is built (trained); the validation set, where different models are evaluated to choose the one performing best in the given setting; and the test set, to assess the error rate of the chosen model in previously unseen cases (27). Validation sets are not always used, especially if the annotated dataset is very small. However, the validation set is critical to find the best combination of the model (e.g., logistic regression) and parameters (e.g., model coefficients). That is done in an iterative process where models with different hyperparameters (e.g., type of regularization or maximal number of decisions in a decision tree) are trained on the training dataset and evaluated on the validation set. This process allows for the algorithm's hyperparameters to be tuned in order to maximize its performance. The best algorithm with the optimal hyperparameters can be trained with both the training and the validation set to obtain the final model to be tested. After this iterative process has converged and the best performing model has been found and trained, the next step is to determine the model's performance dealing with new observations. To do that, the model needs to be tested one final time in an unseen dataset, the test set. The validation set cannot be used because, although the model is not trained on the validation set, there is an information leak from it to the model (27). Every time the model is tuned using the validation set, some information leaks into the model. There are several approaches available to generate these data subsets (11). Usually, the test set is created by leaving out a random sample of the data (frequently around 25%). Then, the training and validation sets can be created, at least by two approaches. The simplest one is splitting the remaining data again into training and validation sets. When the dataset is too small to sacrifice a portion of it to a validation set, k-fold cross-validation is useful. In this method, the sample is split into k (usually, most often arbitrarily, 5, 7, or 10) groups called folds (11). One fold serves as a validation set and the rest as the training set. Analysts repeat this but taking the next fold as the validation set and the rest as the training set. The validation process's final metric is the average of each validation step (k-times). The best model is the one with the highest mean and lowest variation in performance across folds. The test dataset results provide a realistic assessment of the model's error rate and help detect overfitting. Overfitting is a common caveat found when working with small datasets, when the dataset is not representative of the universe, or

when a model is too complex for the available information. One must suspect overfitting when obtaining a good performance on the validation set with poor results on the test set. Overfitting occurs when a model follows too closely the training data, incorporating not only patterns but also anecdotal information to make predictions. Overfitting results in a very accurate prediction of training examples lacking in generalization to adequately predict new, previously unseen cases. Some methods such as penalized regression, random forests, or artificial neural networks have inner processes that include constraints to prevent major overfitting. Working with training and validation sets is also a strategy to detect and handle overfitting.

3.3.2. Model Performance Metrics

Different measurements can evaluate models' performance by comparing models' predictions to the observed data points during the validation or testing evaluation steps. Model performance metrics differ according to the task to solve (i.e., regression or classification). Each metric has strengths and shortcomings (Table 1). In the regression setting, where the outcome is a continuous variable, performance metrics determine a distance between the predicted value and the true response. The most common ones are the mean squared error (with the common acronym MSE), the root MSE (usually abbreviated as RMSE), and mean absolute error (shortened as MAE) (28). While MSE and RMSE are easier to interpret, MAE is more robust to outliers in the observed data. Classification problems with binary or categorical outcomes present more complex scenarios because the idea of distance does not apply. Classifiers can fail by not predicting a positive case, a negative case, or minority data points correctly. As summarized in Table 1, model performance metrics proliferate with the different ways classifiers can fail for classification. Different classification metrics prioritize different aspects of the classification task (28). All of these metrics are calculated using the confusion matrix, a well-known 2-by-2 contingency table in diagnostics research with dichotomous outcomes (Figure 2). Many tasks are associated with standard performance metrics. In machine learning, accuracy is the default performance metric, measuring the ratio of correct predictions. This metric provides a rough but intuitive estimate of overall performance but fails to highlight differences between different kinds of errors. Precision and recall detail different kinds of errors (correct predictions among positive predictions and correct predictions across positive cases, respectively), and the F1 measure is a mean of the two. The Area Under the Curve of Receiver Operating Characteristics (AUC ROC) metric provides a complete assessment of the model because it measures the model's performance at several thresholds. For a more precise definition of these metrics and synonyms in the applied statistics framework, see Table 1 and the Glossary. However, standard metrics often fall short of gaining insight on relevant aspects of the model's performance. Additional metrics may be needed to obtain a more thorough understanding of what may hurt or help performance. The metrics' choice must be in agreement with the study aims. It is possible to select a metric that prioritizes the completeness in the detection of positive cases (i.e., sensitivity, also known as recall) at the expense of a possible increase in false positives, or a metric that emphasizes the correct detection of positive cases even though some cases are not detected (i.e., positive predictive value, also known as precision). For example, Lo-Ciganic et al. (6) tackled the classification problem of predicting the risk of an

	Actually Positive	Actually Negative
Predicted Positive	True Positives (TPs)	False Positives (FPs)
Predicted Negative	False Negatives (FNs)	True Negatives (TNs)

Figure 2. Confusion Matrix. It is the 2-by-2 contingency table with actual outcome categories and the predicted outcome categories by the machine learning model. This table allows us to calculate most of the metrics described in table 1.

opioid overdose among Medicare beneficiaries with at least one opioid prescription. The authors compared algorithms using AUC ROC together with other metrics. While two algorithms (gradient boosting machine and artificial neural networks, see Part I for more details (1)) had the best AUC ROCs, they differed in sensitivity (87% vs. 92%, respectively) and specificity (81% vs. 76%, respectively). The correct interpretation of any given metric is central to the model evaluation process. Researchers need to understand the meaning of the metric and consider a baseline performance to compare the model. This baseline is usually fixed as the performance achieved by predicting each case's outcome by chance or with the least complex machine learning method's performance to address the problem at hand. Thus, several machine learning algorithms can be evaluated by comparing their AUC ROC, as in Lo-Ciganic et al. (6). If only one method is used, its performance can be quantified by comparing it with a random classifier. For example, Whelan and colleagues (29) applied a penalized logistic regression with multiple variables from several neuropsychological domains to identify adolescent binge drinking behavior. After several variable selection steps, the model AUC ROC was higher than 0.5, the AUC ROC value that one would obtain if, instead of a model to predict if an adolescent has binge drinking behavior, one tosses a fair coin. Whether researchers use several trained models or just one, it is critical to determine if the metric performance differences are clinically meaningful. Imbalanced datasets have few cases with the condition of interest over the total number of cases (e.g., the number of patients with an opioid overdose versus other diagnoses in a database from electronic health records (3)). Some model performance metrics are inadequate for imbalanced datasets. For example, in datasets with a high imbalance ratio (e.g., 1 positive per 1000 negative cases), a classifier that systematically assigns the majority class to every case can reach a high accuracy (e.g., 99% in this example). Thus, metrics such as accuracy or AUC ROC can be misleading and may overestimate the model performance (30). The first step to handle this issue is to identify the imbalance ratio in a dataset. The next step is to determine where to focus concerning the study's outcome because each metric prioritizes different performance aspects. There are several ways to work with performance metrics to tackle the imbalance problems. The AUC ROC can be reported together with the imbalance ratio, several performance curves can be analyzed and displayed (e.g., ROC curve, precision-recall curve, and detection error tradeoff curve). A ROC curve region of interest can be selected and used instead of the whole curve. Finally, the available metrics can be customized to a problem, particularly useful for multiclass variable responses (e.g., balanced accuracy or class-weighted metrics) (30,31). In sum, to address the question of how to choose a model given a dataset, it is recommended that researchers evaluate different models, starting by the simplest, prioritize a performance metric according to the problem to solve, assess significance between metrics, and analyze other model aspects such as interpretability, simplicity, (lack of) overfitting, and results' stability, as assessed by variation in cross-validation.

3.4. Communication and reproducibility of results

As Figure 1 displays, the communication of results condenses all previous data analysis stages. Ideally, communication of results is a clear statement interpreting the model(s) selected that address the questions that originated the study in agreement with the study aims and being explicit about assumptions and limitations. Communication of results should include ways so others can understand what, how, and why of the study conducted alongside the findings of one whole iteration of the extracting-knowledge-from-data cycle. A holistic understanding of research results go hand-in-hand with transparency at the reporting stage, that is, not only reporting results but as many as possible elements of how the results were obtained, including, but not limited to sharing the data to train, validate, and test the models and all code used in the analysis cycle reported (12). The reporting of model parameters is mandatory to ensure reproducibility but may be cumbersome, especially if the model has many parameters and possible values, as is the case with more complex models, like neural networks. In those cases, reporting the code solves this issue while also seeing how data were handled at every step. A beneficial side effect of transparency is that it enables other researchers (or your future self) to reproduce previous results and helps uncover mistakes that improve the evolution of knowledge, an ethical obligation when working with addiction data, and efficient use of economic resources. In this sense, as mentioned at the design stage, the EQUATOR Network guidelines are good companions to guide the communication of results (14). We recommend free, open-source tools for the R and Python programming language environments, such as RMarkdown (32) and Jupyter Notebooks (33), for generating reproducible result reports.

4. Discussion

We have presented an encompassing perspective of machine learning implementations. We started from typical caveats that need to be kept under control when machine learning methods are applied. Then we described in more detail a paradigmatic workflow, from experimental design to the communication of results. From the scarce existing addiction research literature using a machine learning approach to addiction data analysis, in the supplementary material, we discussed some examples of real-world addiction problems as illustrations of this workflow's instances, particular issues that emerge, and how to address them. Machine learning techniques need to be applied within an adequate research framework; they cannot be isolated from the whole research process. Therefore, a workflow for machine learning implementation improves the use of the available resources, deals with typical caveats, and reaches sound, descriptively adequate results. Many of the concepts underlying the workflow are very familiar to researchers that work with data, indicating that machine learning is an accessible toolkit, with many points in common to current practices with widespread data-driven research. Thus, the effort-benefit balance is very positive to include these tools in addiction research pipelines. This review leaves many aspects open for the reader to go into further detail. We offered a broad overview of machine learning methods and their implementation to facilitate their acquisition. This two-part review is merely a primer of machine learning for the addiction field, including pointers to valuable resources and references for further reading. We did not expand on specific areas like natural language or image processing because these areas' complex intricacies require lengthier explanations beyond this work scope. The presented workflow is not a rigid and closed structure; it is a flexible framework to incorporate machine learning tools. Regarding addiction research articles, we did not conduct a systematic review; nevertheless, we conducted a comprehensive literature search to include and present relevant examples that illustrate relevant problems in the area. There is no doubt that machine learning can benefit the addiction research field. The difficulties of entering a new area are alleviated by viewing machine learning in a continuum with traditional statistics. We have shown how a good workflow provides a solid basis to reach sound and reliable results. We contribute this two-part review with the hope of facilitating that addiction researchers incorporate these powerful tools into their analytical toolkit.

5. Data availability statement

Data sharing does not apply to this article as no new data were created or analyzed in this study.

6. References

1. Cresta Morgado P, Carusso M, Alonso Alemany L, Acion L. Practical foundations of machine learning for addiction research. Part I. Methods and techniques. *Am J Drug Alcohol Abuse*. Forthcoming.
2. Rajkomar A, Dean J, Kohane I. Machine Learning in Medicine. *N Engl J Med*. 2019 Apr 4;380(14):1347–58.
3. Dong X, Rashidian S, Wang Y, Hajagos J, Kong J, Saltz M, et al. Machine Learning Based Opioid Overdose Prediction Using Electronic Health Records. *AMIA Annu Symp Proc*. 2020 Mar 4;2019:389–98.
4. Hu Z, Jing Y, Xue Y, Fan P, Wang L, Vanyukov M, et al. Analysis of substance use and its outcomes by machine learning: II. Derivation and prediction of the trajectory of substance use severity. *Drug Alcohol Depend*. 2020 Jan;206:1–7.
5. Graves RL, Tufts C, Meisel ZF, Polsky D, Ungar L, Merchant RM. Opioid Discussion in the Twittersphere. *Subst Use Misuse*. 2018 Nov 10;53(13):2132–9.
6. Lo-Ciganic W-H, Huang JL, Zhang HH, Weiss JC, Wu Y, Kwok CK, et al. Evaluation of Machine-Learning Algorithms for Predicting Opioid Overdose Risk Among Medicare Beneficiaries With Opioid Prescriptions. *JAMA Netw Open*. 2019 Mar 22;2(3):e190968.
7. Jing Y, Hu Z, Fan P, Xue Y, Wang L, Tarter RE, et al. Analysis of substance use and its outcomes by machine learning I. Childhood evaluation of liability to substance use disorder. *Drug Alcohol Depend*. 2020 Jan;206:1–6.
8. Sakoglu U, Mete M, Esquivel J, Rubia K, Briggs R, Adinoff B. Classification of cocaine-dependent participants with dynamic functional connectivity from functional magnetic resonance imaging data. *J Neurosci Res*. 2019 Jul;97(7):790–803.
9. Beam AL, Kohane IS. Big Data and Machine Learning in Health Care. *JAMA*. 2018 Apr 3;319(13):1317.
10. Mitchell TM. *Machine Learning*. New York: McGraw-Hill; 1997. 414 p. (McGraw-Hill series in computer science).
11. James G, Witten D, Hastie T, Tibshirani R. *An Introduction to Statistical Learning* [Internet]. New York, NY: Springer New York; 2013 [cited 2020 Jun 10]. 426 p. (Springer Texts in Statistics; vol. 103). Available from: <http://link.springer.com/10.1007/978-1-4614-7138-7>
12. Kass RE, Caffo BS, Davidian M, Meng X-L, Yu B, Reid N. Ten Simple Rules for Effective Statistical Practice. Lewitter F, editor. *PLOS Comput Biol*. 2016 Jun 9;12(6):e1004961.
13. Riley RD, Ensor J, Snell KIE, Harrell FE, Martin GP, Reitsma JB, et al. Calculating the sample size required for developing a clinical prediction model. *BMJ*. 2020 Mar 18;368:m441.
14. EQUATOR network [Internet]. Available from: <https://www.equator-network.org/>
15. Collins GS, Reitsma JB, Altman DG, Moons KGM. Transparent Reporting of a multivariable prediction model for Individual Prognosis Or Diagnosis (TRIPOD): The TRIPOD Statement. *Ann Intern Med*. 2015 Jan 6;162(1):55.
16. Moons KGM, Altman DG, Reitsma JB, Ioannidis JPA, Macaskill P, Steyerberg EW, et al. Transparent Reporting of a multivariable prediction model for Individual Prognosis Or Diagnosis (TRIPOD): Explanation and Elaboration. *Ann Intern Med*. 2015;162:W1–73.
17. Luo W, Phung D, Tran T, Gupta S, Rana S, Karmakar C, et al. Guidelines for Developing and Reporting Machine Learning Predictive Models in Biomedical Research: A Multidisciplinary View. *J Med INTERNET Res*. 18(12):e323.
18. Liu X, Rivera SC, Moher D, Calvert MJ, Denniston AK, SPIRIT-AI, et al. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the CONSORT-AI extension. *Nat Med*. 2020;26:11.

19. Collins GS, Moons KGM. Reporting of artificial intelligence prediction models. *The Lancet*. 2019 Apr;393(10181):1577–9.
20. Sounderajah V, Ashrafian H, Aggarwal R, De Fauw J, Denniston AK, Greaves F, et al. Developing specific reporting guidelines for diagnostic accuracy studies assessing AI interventions: The STARD-AI Steering Group. *Nat Med*. 2020 Jun;26(6):807–8.
21. Bossuyt PM, Reitsma JB, Bruns DE, Gatsonis CA, Glasziou PP, Irwig L, et al. STARD 2015: an updated list of essential items for reporting diagnostic accuracy studies. *BMJ*. 2015 Oct 28;351:h5527.
22. van Smeden M, Penning de Vries BBL, Nab L, Groenwold RHH. Approaches to Addressing Missing Values, Measurement Error and Confounding in Epidemiologic Studies. *J Clin Epidemiol* [Internet]. 2020 Nov [cited 2020 Nov 30]; Available from: <https://linkinghub.elsevier.com/retrieve/pii/S0895435620311756>
23. Andreas C. M, Guido S. Introduction to Machine Learning with Python. First. California, USA: O'Reilly Media, Inc; 2017.
24. Duboue P. The Art of Feature Engineering: Essentials for Machine Learning [Internet]. 1st ed. Cambridge University Press; 2020 [cited 2020 Nov 25]. Available from: <https://www.cambridge.org/core/product/identifier/9781108671682/type/book>
25. Peng RD, Matsui E. The Art of Data Science [Internet]. 2017 [cited 2020 Oct 28]. Available from: <https://bookdown.org/rdpeng/artofdatascience/>
26. The Turing Way Community, Becky A, Bowler L, Gibson S, Herterich P, Higman R, et al. The Turing Way: A Handbook for Reproducible Data Science (Version v0.0.4) [Internet]. Zenodo. 2019 [cited 2020 Dec 1]. Available from: <http://doi.org/10.5281/zenodo.3233986>
27. Chollet F, Allaire JJ. Deep learning with R. Shelter Island, NY: Manning Publications Co; 2018. 335 p.
28. Géron A. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow. 2nd ed. California, USA: O'Reilly Media, Inc; 2019. 851 p.
29. Whelan R, Watts R, Orr CA, Althoff RR, Artiges E, Banaschewski T, et al. Neuropsychosocial profiles of current and future adolescent alcohol misusers. *Nature*. 2014 Aug;512(7513):185–9.
30. Brabec J, Machlica L. Bad practices in evaluation methodology relevant to class-imbalanced problems. In Montréal, Canada.; 2018 [cited 2020 Nov 30]. Available from: <http://arxiv.org/abs/1812.01388>
31. Gupta A, Tatbul N, Marcus R, Zhou S, Lee I, Gottschlich J. Class-Weighted Evaluation Metrics for Imbalanced Data Classification. *ArXiv201005995 Cs* [Internet]. 2020 Oct 12 [cited 2020 Nov 30]; Available from: <http://arxiv.org/abs/2010.05995>
32. Wickham H, Golemund G. R for Data Science [Internet]. First. California, USA: O'Reilly Media, Inc; 2017. Available from: <https://r4ds.had.co.nz/>
33. Kluyver T, Ragan-Kelley B, Pérez F, Granger B, Bussonnier M, Frederic J, et al. Jupyter Notebooks—a publishing format for reproducible computational workflows. In IOS Press; 2016. p. 87–90.

Table 1. Machine learning performance measurements with their counterpart in the statistics, strengths, and shortcomings.

Term	Description	Strengths	Shortcomings
Continuous response variable			
MAE (Mean Absolute Error)	It is the average of the errors. The errors are the distance between the predicted value and the true response.	It is robust to outliers.	It is mathematically untractable and computationally more challenging.
MSE (Mean Squared Error)	It is the average of the squares of the errors.	It is commonly used and is preferred when outliers are rare. Mathematically tractable.	It is sensitive to outliers.
RMSE (Root Mean Squared Error)	It is the root of the MSE.	Easy to interpret because it is in the same units as the mean.	It is sensitive to outliers.
Categorical response variable			
Accuracy $(TP + TN)/\text{Total Population}$	It is the percentage of correct predictions from the total number of cases. This metric provides a rough but intuitive estimate of overall performance.	Easy to interpret. It considers both positive and negative predictions.	It can overestimate performance with an imbalanced dataset. It does not acknowledge the differences between different kinds of errors.
Precision. Synonymous: Positive Predictive Value $TP/(TP + FP)$	It is the proportion of positive predictions that are true positives. It is the ratio of true positive predictions to the total number of positive predictions. It accounts for the model's reliability for a positive prediction.	It is a good measure of the reliability of positive predictions (penalizes false detections). It works for imbalanced datasets.	It does not penalize false negatives. A high precision may be trivial with a large number of false negatives.
Negative Predictive Value $TN/(TN + FN)$	It is the proportion of negative predictions that are true negatives. It is calculated as the ratio of true negative predictions to the total number of negative predictions.	It shows the reliability over negative predictions. It could complement other metrics such as precision. It works for imbalanced datasets.	It is less used than positive predictive value (precision).
Recall. Synonymous: Sensitivity. $TP/(TP + FN)$	It is the proportion of positive cases correctly identified. It is calculated as the ratio of true positive predictions to the total number of actual positive cases. This accounts for the model's coverage of positive.	It is a good measure when the primary focus is the identification of as many cases as possible. It works for imbalanced datasets.	It does not penalize false positives (i.e. may allow for low reliability of positive predictions)
Specificity $TN/(TN + FP)$	It is the proportion of negative cases correctly identified. It is the ratio of true negative predictions to the total number of actual negative cases.	It is easy to interpret and commonly used. A high value ensures that a negative case predicted is probably true.	As it places the focus only on negative cases, it requires a complementary measure; usually, sensitivity.
False Positive Rate. Synonyms: $1 - \text{Specificity}$, Type I error. $FP/(TN + FP)$	It is the proportion of negative cases incorrectly classified. It is the ratio of false positive predictions to the total number of actual negative cases.	It penalizes false positives considering the number of true negatives which could be useful in imbalanced datasets.	It requires being shown with another metric to complement it.

Table 1. (Continued.)

Positive likelihood ratio $\text{Sensitivity}/(1-\text{Specificity})$	It is a ratio between the probability of having a true positive over the probability of having a false positive. It is calculated by dividing specificity (recall) by 1-specificity (false positive rate.). Unrelated to the likelihood-ratio test.	It summarizes two probabilities. When the ratio is greater than 1 the method is useful and reliable.	It is not frequently used.
Negative likelihood ratio $(1-\text{Specificity})/\text{Sensitivity}$	It is a ratio between the probability of having a false negative over the probability of having a true negative. It is calculated by dividing 1-specificity (false positive rate) by specificity (recall). Unrelated to the likelihood-ratio test.	It summarizes two probabilities. When the ratio is lower than 1 the method is useful and reliable.	It is not frequently used.
Number needed to treat or to evaluate	It is the number of individuals you need to treat (or test) to prevent one additional bad outcome. It is calculated by: $\text{NTT} = 1/(\text{event rate in control group} - \text{event rate in experimental group})$.	It is informative and shows the impact (as a number of individuals who need to be treated) of the intervention that is under study.	It is not frequently used as a performance measure to train and evaluate machine learning models.
Area Under the Curve of Receiver Operating Characteristics (AUC ROC)	The Receiver Operating Characteristics curve plots the probability of detecting true signal (sensitivity) and false signal (1 - specificity) for an entire range of possible cutpoints. The AUC ROC measures the model's performance while simultaneously considering several thresholds.	The AUC ROC considers all decision thresholds, thus resulting in a more complete assessment of the model. There are techniques to find the best threshold. It combines both sensitivity (recall) and False Positive Rate.	It may be difficult to interpret. It can overestimate performance with an unbalanced dataset.
F1 score	It is the harmonic mean of the precision and the recall. The score is high if precision and recall are also high.	It is a way to summarize precision and recall in one measure. It considers coverage and reliability.	It is difficult to interpret. Both precision and recall may require different considerations, though the metric may be adjusted in that case.