Machine Learning

Jennie E. Brand, Bernard Koch, Jiahui Xu

PWP-CCPR-2019-001

June 10, 2019

California Center for Population Research
On-Line Working Paper Series

Machine Learning*

Jennie E. Brand† Bernard Koch‡ Jiahui Xu§

June 10, 2019

Abstract

Machine learning is a statistical and computational approach to extracting important patterns and trends in data. This entry is an overview of machine learning methods for social science research. Supervised learning methods are discussed, including generalized linear models, support vector machines, *k*-nearest neighbor, artificial neural networks and deep learning, decision trees, and ensemble methods based on decision trees. Several important considerations relevant to supervised learning algorithms are noted, including the use of training and test data and cross-validation, loss optimization and evaluation metrics, bias-variance tradeoff, and overfitting and regularization strategies. Unsupervised learning methods are also discussed, including *k*-means clustering, hierarchical clustering, network community detection, principal component analysis, and *t*-distributed stochastic neighbor embedding. A section on text analysis incorporates supervised and unsupervised learning of documents and neural networks. New developments at the intersection of machine learning methods and causal inference are discussed. Key limitations and considerations for adopting these methods in empirical social science research concludes the entry.

---

* Authorship is listed in alphabetical order. The first author benefited from facilities and resources provided by the California Center for Population Research at UCLA (CCPR), which receives core support (P2C-HD041022) from the Eunice Kennedy Shriver National Institute of Child Health and Human Development (NICHD). The ideas expressed herein are those of the authors.

† Professor of Sociology and Statistics, UCLA, and Director, California Center for Population Research, and Co-Director, Center for Social Statistics, brand@soc.ucla.edu.

‡ Sociology Masters Student, UCLA, bernardkoch@g.ucla.edu.

§ Masters in Applied Economics, UCLA, jiahuixu@ucla.edu.

SAGE Research Methods Foundations:

Submission for consideration

Metadata Sheet

Entry Title		Machine Learning
Authors: [Note: Authors are listed alphabetically.]		
1	Name	Jennie E. Brand
	Affiliation, country	University of California, Los Angeles, USA
	Lead author email	brand@soc.ucla.edu
	SAGE Author ID	[office use only]
2	Name	Bernard Koch
	Affiliation, country	University of California, Los Angeles, USA
	SAGE Author ID	[office use only]
3	Name	Jiahui Xu
	Affiliation, country	University of California, Los Angeles, USA
	SAGE Author ID	[office use only]

All author bio(s)	Jennie E. Brand is Professor of Sociology and Statistics at UCLA. She is Director of the California Center for Population Research and Co-Director of the Center for Social Statistics at UCLA. She is Chair of the Methodology Section of the American Sociological Association, and a member of the Board of Overseers of the General Social Survey (GSS), and a member of the Technical Review Committee for the National Longitudinal Surveys Program at the Bureau of Labor Statistics. Her research centers on inequality and its implications for various outcomes that indicate life chances, and a methodological focus on causal inference and the application and innovation of statistical models for panel data. Current research projects include evaluation of heterogeneity in the effects of education on socioeconomic outcomes and the social-psychological and socioeconomic consequences of disruptive events. Bernard Koch is a Ph.D. student in the Sociology Department at the University of California-Los Angeles. He is interested in the science of science, cultural evolution, and computational methods. For his dissertation, he'd like to focus on applications of deep learning to causal inference and network problems to understand how we can make science more efficient, productive, and equitable. Jiahui Xu is a master's student in Applied Economics at the University of California - Los Angeles. Her research interests lie in machine learning, causal inference and demography. She currently works on applying machine learning to causal inference and evaluation of sociological effect heterogeneity.

Section to be filled in by authors:

Discipline(s)	Economics [D16] Sociology [D1]

	Public Policy and Social Policy [D6] Education [D2] Computer Science [D20]
Method Group	Quantitative
Entry Category	Flagship
Entry Size (approximate word count)	10,000
Key Figures mentioned	
Key Organizations mentioned	
Abstract	Machine learning is a statistical and computational approach to extracting important patterns and trends in data. This entry is an overview of machine learning methods for social science research. Supervised learning methods are discussed, including generalized linear models, support vector machines, k-nearest neighbor, artificial neural networks and deep learning, decision trees, and ensemble methods based on decision trees. Several important considerations relevant to supervised learning algorithms are noted, including the use of training and test data and cross-validation, loss optimization and evaluation metrics, bias-variance tradeoff, and overfitting and regularization strategies. Unsupervised learning methods are also discussed, including k-means clustering, hierarchical clustering, network community detection, principal component

	analysis, and *t*-distributed stochastic neighbor embedding. A section on text analysis incorporates supervised and unsupervised learning of documents and neural networks. New developments at the intersection of machine learning methods and causal inference are discussed. Key limitations and considerations for adopting these methods in empirical social science research concludes the entry.

New identifiers for this case:	
Copyright year	2020
Copyright statement	© SAGE Publications Ltd 2020
DOI	10.4135/
URL	http://methods.sagepub.com/foundations/
URI	

Machine Learning

Introduction

Machine learning is a computational and statistical approach to extracting patterns and trends from data (Maini and Sabri 2017). A machine learning algorithm is also defined as a computer program that learns from experience with data with respect to some class of tasks, and improves its performance with greater experience (Mitchell 1998). Machine learning has widespread applications to many scientific fields and has gained a considerable amount of attention and use in recent years in the era of high-speed computing and big data. As the tasks of statistical sciences have expanded in scope and complexity, machine learning has quickly permeated many scientific fields, from statistics and computer science to the social sciences.

There has been some debate, particularly in statistics and the social sciences, about the degree of difference between machine learning and classical statistics. Many describe the field as statistical learning. Machine learning methods overlap with methods that have been used in applied statistics and social sciences for many decades, although it is only recently that these fields have adopted machine learning methods in earnest. At the same

time, statisticians and social scientists have made progress in bridging machine learning methods to areas central to statistics and social science, like causal inference (Athey 2018).

Machine learning algorithms can be categorized into supervised, unsupervised, and reinforcement learning. Supervised learning algorithms learn from training data that contains both inputs (also called, independent variables, covariates, and inputs) and labeled outcomes (also called dependent variables, outcomes, and responses). Supervised learning tasks where the outcome variable is categorical are typically referred to as classification tasks; tasks where the outcome is continuous are called regression tasks. Some common supervised learning approaches included generalized linear models (e.g. logistic regression), kernel methods (e.g., support vector machines), decision tree-based methods, and artificial neural networks (e.g., deep learning).

In contrast to supervised learning, unsupervised learning algorithms build models with data only on the features, without labels on the outcomes. These algorithms are usually used for the purpose of clustering (e.g., discovering smaller groups within the data) or dimensionality reduction (e.g., reducing the number of features/covariates to a smaller, more informative number of latent features). These algorithms build models with data only on the input features, without labels on the outputs. Some common unsupervised learning techniques include k-means clustering, hierarchical clustering, principle components analysis, t-distributed stochastic neighbor embedding (t-SNE). These techniques may be an

intermediate data analysis step in the construction of covariates or outcome variables for supervised learning.

While supervised learning uses data with outcomes (i.e., labels are provided), and unsupervised learning does not (i.e., no labels are provided), an additional approach, reinforcing learning, lies in the middle. Reinforcement learning is concerned with how agents should take actions in environments to maximize reward. As supervised learning has seen the most widespread use in the social sciences, this entry will devote considerable time to supervised learning, less time to unsupervised learning, and no additional time to reinforcement learning.

A Brief History of Machine Learning: From Computer Science to Computational Social Science

Machine learning is a subset of artificial intelligence (AI). Since Alan Turing first developed the concept of the Turing Test in 1950, researchers in computer science have been drawn to the study of AI. Turing's paper considered the question of whether machines can "think." Following the Dartmouth conferences of 1956, known as the "birth of AI"

(Crevier 1993), Arthur Samuel coined the term "machine learning" to describe the way in which machines learn from data. Machine learning methods were later found to be similar to methods in statistics. Many of these early statistics models ultimately fell outside of the scope of AI and computer science. Machine learning grew in the 1990s as a means by which to apply statistics and probability theory to effectively utilize increasing digitized information.

The term "data science" unifies the fields of machine learning, statistics, and data analysis. Data science is a multi-disciplinary field, drawing on mathematics, statistics, computer and information science, that uses scientific methods and algorithms to generate knowledge from data. The term was used by Peter Naur in 1960 to refer to computer science, but it had not become widely used until recently. In recent years, data science has become a widely used phrase and "data scientist" a coveted, if poorly defined, position. Many universities are either instituting or rebranding data science divisions, departments, centers, and labs. Another emerging phrase in the social sciences is "computational social science." Computational social science is data science applied to social science applications, with an emphasis on network analysis, use of digital trace data, and leveraging online experiments and surveys.

The transformative impact of machine learning on social science research and policy is well underway. Machine learning gives social science researchers tools to leverage the

large, rich, and often unstructured data that have become increasingly available. While these data benefit from being large relative to traditional social science surveys, they come with the own unique set of challenges, such as population representativeness and quality of measures. With progress in software and hardware, as well as statistical approaches, machine learning methods are a powerful tool for social scientists to leverage these new and existing data. This does not, however, negate the importance of human interpretative ability and social theory in social science research (Pearl 2018).

Supervised Learning

Supervised learning algorithms learn to identify labeled outcomes from training data, and then can be used to predict the outcomes of new unseen test data. The algorithm observes both the covariates and the outcomes. For example, to teach a supervised algorithm how to distinguish between democrats or republicans, the researcher would provide a training sample of individuals with a set of covariates and an outcome labeled as democrat and republican. After training, the algorithm is used to predict the party affiliations of individuals in a new dataset. As noted above, tasks where the outcome variable is categorical are referred to as classification tasks, and tasks where the outcome is continuous are called regression tasks. Many algorithms can be adapted for both regression and classification tasks. For example, a researcher might use logistic regression to predict the probability that

individuals attend college (regression) or round this probability at a decision boundary to predict whether or not individuals attend college (classification).

There are a wide range of supervised learning approaches and no single algorithm performs best in all applications. Some common supervised learning approaches are: (a) generalized linear models, (e.g., OLS and logistic regression); (b) kernel methods (e.g. support vector machines), (c) k-nearest neighbors, (d) decision tree methods (e.g. classification and regression trees (CART), random forests, and gradient boosting machines), and (e) artificial neural networks (i.e., deep learning). The rest of this section proceeds follows. First, we describe the general workflow of supervised learning with a focus first on regression tasks. Second, we explain how supervised regression algorithms are trained and introduce the bias-variance tradeoff, as well as the concept of regularization. Finally, we discuss classification tasks and the precision-recall tradeoff. We then briefly introduce readers to the different classes of supervised learning algorithms.

Training, Cross-validation, and Testing Data

The first step in a machine learning analysis is to divide the data into a labeled training set for optimizing the model parameters, and a testing set of unlabeled data that we wish to predict outcomes for. To avoid overfitting, the testing set should not be touched until the final parameters of the model have been chosen and this set will be used only once.

Researchers reserve a portion of the training set called the cross-validation set to evaluate how well the algorithm can perform out-of-sample prediction before showing it the testing data, and to tune hyperparameters (i.e., model parameters that need to be adjusted by the researcher) that can improve performance. Typically, researchers reserve between 10-50% of their training data for cross-validation, depending on the size of their dataset. However, techniques such as k-fold cross-validation can reduce the amount of cross-validation data needed. In k-fold cross-validation, researchers divide the data into k groups and build a model on k-1 groups while the remaining group is used for validating. This subsampling process is repeated until all the k groups have been used as the test set. Figure 1 is an example of k-fold cross-validation. The sample is partitioned into a training set and a testing set, and the training set is subsequently partitioned into a cross-validation set to tune hyperparameters. To sample different hyperparameter values during cross-validation, researchers often use grid search (i.e., explore all possibilities) or random search (i.e. explore random possibilities), although more sophisticated strategies such Bayesian optimization are also possible.

Loss Optimization and Evaluation Metrics

During training, supervised learning algorithms optimize in-sample performance with respect to a loss (sometimes called objective or cost) function. After training, researchers

use evaluation metrics to assess out-of-sample predictive performance of the model. Although most supervised learning algorithms can be used for both regression (continuous outcome) and classification (categorical outcome), regression and classification tasks typically have different loss optimization functions and evaluation metrics. Moreover, although the same metrics are often used for in-sample loss and out-of-sample evaluation in regression, researchers typically use different metrics for training and evaluation in classification tasks. This section will focus first on functions for loss and evaluation in regression, and then on classification.

Regression Loss and Evaluation Metrics

The most common loss function for regression is the mean square error (*MSE*), or quadratic loss function. The *MSE* measures the mean of the square of the errors between the predicted value and the actual value:

$$MSE = \frac{1}{n}\sum_{i=1}^{n}(y_i - \left(\beta_0 + \sum_{j=1}^{p}\beta_j\, x_{ij}\right))^2.$$

Sometimes the square root of this function, the root mean squared error (RMSE), is used. The *MSE* is always non-negative, and values closer to zero indicate better model accuracy. In regression training, the algorithm is trained to minimize *MSE* in the training data; researchers are most interested however, in how well the model performs on the test data.

The *MSE* is also the most common evaluation metric. There is no guarantee that the model with the smallest training set *MSE* will have the smallest test set *MSE*.

Mean absolute error (*MAE*) is an alternative regression loss and evaluation function that is more robust to outliers. *MAE* is the sum of absolute differences between the predicted and actual value rather than the sum of squared differences.

$$MAE = \frac{1}{n}\sum_{i=1}^{n} |y_i - \left(\beta_0 + \sum_{j=1}^{p} \beta_j \, x_{ij}\right)|.$$

Although the *MSE* is by far the most common loss/evaluation metric, there may be specific tasks where a more specialized function like Huber loss, log-cosh, or quantile loss is appropriate. Beyond these functions, many of the model selection metrics that social scientists are already familiar with, such as R^2, are also useful for evaluating out-of-sample performance.

Classification Loss and Evaluation Metrics

The most common classification loss function is the log loss. Rather than optimizing just the accuracy of classification, log-loss optimizes the model's probabilistic estimates of which class the data belong to. Suppose there are two possible class labels, y is a binary variable (0/1) indicating whether class label c is the correct classification for a unit i, and p is the model's predicted probability that unit i is class c, the log loss function is:

$$-(y \times log(p) + (1 - y)log(1 - p)).$$

A perfectly predictive model would have a log loss of 0, and the log loss increases as the predicted probability deviates from the actual label. For example, if the true label had a value of 1.0, the log loss function slowly decreases as the predicted probability nears 1.0. Log-loss can be extended to multiclass problems by summing the first term in the equation above for each class. Yet not all classifiers extend generically to multi-classification. Focal loss, KL divergence or relative entropy, exponential loss, and hinge loss (particularly for SVMs) are other possible loss functions for classification.

Classification accuracy (the percentage of correct classifications) is a simple metric to evaluate classification algorithms, but it does not completely capture the performance of a classifier. Research instead evaluate the performance of a classifier using precision and recall. Precision is defined as:

$$P = True\ Positives(TP)\ /\ True\ Positives + False\ Positives\ (FP),$$

and recall is defined as:

$$R = True\ Positives\ /\ True\ Positives + False\ Negatives\ (FN),$$

If a classifier has high precision, it will rarely misclassify a false positive as true positive. If a classifier has high recall, it will rarely miss any positives that should be true positives. Generally, there is a tradeoff between precision-recall, and there may be times where a researcher favors a more conservative, high-precision or a more sensitive, high-recall

classifier. In binary classification curves, researchers can demonstrate the precision-recall tradeoff using a receiver-operating curve (ROC) plot. Ideally, researchers want a classifier that minimizes the tradeoff. The area under the ROC curve (AUC) is often used as a reductive statistic for comparing classifiers. Yet ROC curves do not extend easily to multiclassification problems. In this case, researchers may turn to the confusion matrix, which displays the number of each true label categorized into each class. Figure 2 displays ROC and the confusion metric. Another approach to summarizing the precision-recall tradeoff is the $F1$ score:

$$F1 = 2\left[\frac{1}{\frac{1}{P}+\frac{1}{R}}\right].$$

Unlike ROC/AUC, $F1$ can be easily extended to multiclassification problems.

Bias-Variance Tradeoff

In all supervised machine learning tasks, there is a fundamental tradeoff between a model's ability to achieve the optimal in-sample fit and its ability to generalize to new, out-of-sample data. In statistical learning theory, this is known as the bias-variance tradeoff. Imagine you have a single datum with covariates x_i and outcome y_i and you place it in many different datasets, and train a supervised learning model *f(x)* on each of these datasets. Error due to bias is the difference between the average (expected) predicted outcome y

across all models/datasets and the actual value of y. That is, a learning algorithm is biased if it systematically incorrectly predicts y. In other words, the bias is the difference between the true values and the algorithm's prediction (in expectation) when trained on multiple datasets. Error due to variance describes the variability of the predictions of y across all models/datasets.

Recall the equation for the *MSE*:

$$MSE = E[(y_i - \widehat{f(x)})^2].$$

The equation for the *MSE* can be decomposed into errors in prediction as errors arising from bias and those from variance as follows:

$$MSE = \left(E\left[\widehat{f(x)}\right] - f(x)\right)^2 + E\left[\left(\widehat{f(x)} - E[\widehat{f(x)}]\right)^2\right] + \sigma_\varepsilon^2.$$

In other words, the *MSE* can be written as bias^2 + variance + irreducible error (Hastie et al. 2009). The irreducible error is the noise in the true relationship that cannot be reduced by any model. The prediction error of a learned classifier is related to the sum of the bias and the variance of the learning algorithm, i.e. a bias-variance tradeoff (Geman, Bienenstock, and Doursat 1992).

Ideally, a model should be neither overly complex nor overly simple. Generally, a more flexible, complex model decreases bias but increases variance. A less complex model essentially ignores relevant information, and error due to bias increases while error due to variance decreases. Figure 3 displays this relationship. A learning algorithm must be flexible

enough to fit the data and minimize bias, while simultaneously not so flexible that variance is high and each data set is fit differently. If a model exceeding the line of optimum complexity is selected, it would be overfitting the data; otherwise, it is underfit. Despite a preconceived leaning toward minimizing bias, researchers thus prefer an estimator with some bias to reduce variance and minimize the loss function.

The size of training data and the number of features influence model building capacity. If the true relationship between features and the outcome is simple, then a small amount of training data and an inflexible learning algorithm with high bias and low variance will be sufficient to learn that function. If the true function is complex, more data and a more flexible algorithm with low bias and high variance is needed to approximate the true function. Likewise, if the feature or X space has high dimensionality, the learning may be difficult. This can be true even if the true function only depends on a small subset of those features.

Overfitting and Regularization

Overfitting occurs when the supervised learning model learns too exactly the particularities of the training data, and may therefore fail to predict future observations reliably in the test data, or more generally to future data. An overfit model contains more parameters than can be justified given the data, and extracts noise from the data as if that

noise or variation represented true model structure. An overfit model has high variance and low bias. Figure 4 displays an example of overfitting an OLS regression model. The other alternative, of course, is that the model is underfit. This occurs with a simple model which has low variance and high bias.

The best strategy to prevent overfitting Is to increase the size of training data. Machine learning algorithms clearly perform better when more data are available to build the model. But increasing data size is not always feasible. Instead, regularization encompasses a range of strategies to prevent overfitting, improve generalization, and reduce model complexity. There are many different approaches to regularization, including reducing the number of covariates using dimensionality reduction, feature selection algorithms, or domain knowledge.

Shrinkage methods reduce the magnitude of parameters by applying a constraint and penalty on the complexity of the model. The two most common shrinkage penalties are L1 or Least Absolute Shrinkage and Selection Operator (LASSO) regression (Tibshirani 1996) and L2, Tikhonov, or ridge regularization. Although LASSO and ridge penalties can be applied to any continuous loss function, this section mainly focuses on their application to generalized linear models (GLMS) (e.g. OLS, logistic regression). For researchers using GLMs, the key difference between LASSO and ridge regression is that LASSO selects a subset of covariates for the final model, and effectively shrinks parameters that do not

contribute to the predictivity of the model to zero, whereas ridge regression shrinks large parameters but does not allow the coefficients to be zero. LASSO therefore results in a simpler, more interpretable model.

Consider a sample of N units, one outcome y_i, p covariates, and $x_i = (x_1, x_2, ..., x_p)^T$ is the covariate vector for the i^{th} case. The objective of LASSO is to solve:

$$\min_{\beta_0,\beta}\left\{\frac{1}{N}\sum_{i=1}^{N}\left(y_i-\left(\beta_0+\sum_{j=1}^{p}\beta_j\,x_{ij}\right)\right)^2\right\},$$

subject to the constraint

$$\sum_{j=1}^{p}|\beta_j|\le t,$$

where t is a prespecified free parameter that determines the degree of regularization. The function with the LASSO or L1 regularization can then be written as:

$$L1=\sum_{i=1}^{N}\left(y_i-\left(\beta_0+\sum_{j=1}^{p}\beta_j\,x_{ij}\right)\right)^2+\lambda\sum_{j=1}^{p}|\beta_j|,$$

where the last term is the penalty, or regularization, term. The λ term is a hyperparameter, which can be tuned in order to improve performance. The ridge regression or L2 regularization is associated with a different penalty term:

$$L2=\sum_{i=1}^{N}\left(y_i-\left(\beta_0+\sum_{j=1}^{p}\beta_j\,x_{ij}\right)\right)^2+\lambda\sum_{j=1}^{p}\beta_j^2,$$

where the last term is the penalty term added. L1 and L2 regularization both prevent overfitting by shrinking the parameters as the penalty term λ increases. If λ is zero, both forms revert to OLS regression. We are faced with a bias-variance tradeoff in choosing the value of λ. As λ increases, the estimators increase bias and minimize variance. As we can see from the different penalty terms for LASSO and ridge, the ridge regularization includes a squared term, and thus the coefficient will never be zero.

Supervised Learning Algorithms

Generalized Linear Models

Generalized linear models, particularly linear and logistic regression, are among the simplest and most widely used supervised learning algorithms. Despite its name, logistic regression is mainly used for binary classification tasks by applying a threshold to the probability generated by the model (probabilities greater than the threshold are assigned to one class, and below the threshold to the other). An example usage of logistic regression might be to classify tumors as malignant or benign based on observed characteristics of the tumor.

In a logistic regression algorithm, the objective is to classify two classes, i.e. y $\in\{0, 1\}$. Log-odds for the value labeled 1 is a linear combination of one or more discrete or covariates. The log odds can also be understood as the log of the probability the event will

occur (or the output is 1) divided by the probability that it does not occur (or the output is 0):

$$logit(p) = log\frac{p}{1-p}$$

This is a linear model, as the log-odds are a linear combination of the inputs. The covariates enter the logistic regression by passing through a sigmoid (*S*-shaped) function. The logistic function is given by the inverse logit.

Consider the tumor example in more detail. Suppose researchers have the features of tumors (including tumor size, position, etc.). The purpose is to use these features to predict tumors, which can be interpreted as probabilities from the logistic function. If the probability for one tumor is greater than 0.5 (where $x = x^*$ when the probability is 0.5), researchers can predict that the tumor is malignant. Otherwise, it is assumed to be benign. Researchers might call 0.5 the decision rule. The vertical line $x = x^*$ is assumed to be a decision boundary which separates the tumors (y=1) from the non-tumors (y=0). Clearly, the data may not, however, be separable at this decision boundary. This example could be extended to incorporate high-dimensional features as well as multiclass classification, such as filtering emails to various categories.

Support Vector Machines

Like logistic regression, support vector machines (SVMs) learn a binary decision boundary between two classes. However, they use a geometric approach rather than probabilistic approach. If the feature space has p dimensions, SVMs will learn a *p-1*-dimensional hyperplane (a two-dimensional hyperplane is a line, a three-dimensional hyperplane is a plane, and so on) to divide the two classes, such that this hyperplane is maximally distant from the nearest training points in each class (i.e., maximizes the "margin" between the hyperplane and each class). Finding the hyperplane that meets the criteria is only computationally feasible by using the "kernel trick," which allows the algorithm to compute dot products (i.e., distances) between points in high-dimensional spaces simply by transforming them using a kernel function (James et al. 2017). Even with p-1 dimensions of flexibility, this "hard" margin approach can lead to overfitting of outliers so "soft margin" SVMs allow a few training units to be misclassified, thus decreasing variance.

k-Nearest Neighbor

The k-nearest neighbor (kNN) is a simple, non-parametric algorithm for both classification and regression. When used for classification, the output is class membership; when used for regression, the output is the value of the response. There is essentially no training phase in kNN. In testing, a new datum is assigned to the class of the plurality (in classification) or mean outcome (in regression) of k nearest training neighbors of the datum

with respect to some distance metric (usually Euclidean, but others such as Hamming are also common). In classification, if k=3 then the new datum is assigned to the majority class of the three nearest neighbors. Higher values of k reduce overfitting, but when k is too high the boundaries between classes blur. To address this, researchers can assign weights to the neighbors, such that nearer neighbors are assigned higher weights. A common weighting technique is to assign each neighbor a weight of $1/d$, where d is the distance to the neighbor.

Artificial Neural Networks and Deep Learning

Artificial neural networks are sophisticated supervised learning algorithms inspired by the human brain. Much like neurons in the brain firing due to impulses from other connected neurons, each node in a neural network takes a weighted combination of the outputs of incoming neurons as its input. This weighted combination is then transformed by an activation function that outputs a single scalar to any neurons to which it is connected. The simplest neural network is called a perceptron where input nodes for each feature feed to a single output node. When the activation function for this neuron is a sigmoid (i.e., logistic) function, this model corresponds exactly with logistic regression. However, the power of neural networks comes from the additional layers of neurons between the input and output layer to create multilayer perceptron (MLPs). With a single fully-connected "hidden" layer of neurons, a neural network can theoretically approximate any continuous

convex function. Neural networks are trained by "backpropagating" the final loss to the weight's in previous layers using calculus's chain rule. Figure 5 is an example of a multi-layer perceptron.

Despite this flexibility for modeling non-linear functions, neural networks have lagged behind other supervised learning approaches in performance for decades because of the computational expense and large amount of data needed to train these models. Starting In around 2011, neural networks with novel architectures began to outperform other supervised learning approaches in what is now collectively called “deep learning.” Two of the most enduring and influential ideas in deep learning are the concepts of convolution and recurrence. Convolutional neural networks are regularized feed forward networks that force neurons to hierarchically learn the complexity of the feature space using a series of convolution and pooling layers. Unlike feedforward networks which are directed and acyclic, recurrent networks have neurons that do not just feed forward, but also connect back to earlier layers in the network. Because they give the network a type of memory, recurrent architectures are popular in natural language processing.

MLPs are not widely used by social science researchers as classifiers/regressors because: (1) unlike decision trees or generalized linear models, they are black-box algorithms with uninterpretable parameters, and (2) they require large amounts of data to train correctly. However, deep learning will become increasingly common in social science as

researchers build familiarity with pre-trained deep learning models, gain access to larger and more complex datasets, and deep learning integrates with causal Inference.

Decision Trees

Decision trees (DTs) are a widely used non-parametric learner for both classification and regression that recursively split the data into increasingly smaller subsets where datapoints are more similar (i.e., have a smaller variance in regression; are more homogenous in classification). The resulting hierarchical data structure can be presented by a tree. These "white-box" models are becoming popular in social science because they are simple to understand and interpret. The most popular DT algorithm is classification and regression trees (CART) (Breiman 1984).

During training, DTs are built from the "root" by recursively: (1) selecting a feature, (2) selecting a threshold/cutoff point on that feature that makes the remaining data on both sides of the threshold more similar, and (3) partitioning the remaining data into two smaller subsets on that feature and threshold. In theory, these steps are repeated until the terminal subsets (i.e., the "leaves") contain only a single data point. In practice, researchers use regularization techniques so that each leaf contains a group of similar training points. In regression, each leaf represents the average value of y for data in that leaf. In classification, each leaf represents the plurality class for the data in that leaf. During testing, researchers

can apply the sequence of decisions used to construct the tree to place each testing datum into the most appropriate leaf.

DTs determine which feature and threshold to split the data on by minimizing an in-sample loss function. Because it is impossible to explore all possible sequences of decisions, DTs minimize this loss function only on the current subset of data at each split, rather than on the entire dataset. This "greedy" approach is efficient, but may result in a locally optimal rather than globally optimal solution. At every split, the loss function tries to maximize homogeneity within leaves. In so doing, researchers maximize entropy across leaves. For regression problems, the loss function minimizes variance with the leaf:

$$c = \sum (Y_i - \bar{Y})^2$$

where $\bar{Y}$ is the mean of Y. For classification problems, there are several possible loss functions: misclassification rate, entropy (deviance), and the Gini impurity. For example, the Gini impurity cost function is an indicator of how pure the two groups will be with respect to class heterogeneity after the split. Formally, the function can be written as:

$$g = \sum_{c=1}^{J} p_c(1 - p_c) = 1 - \sum_{c=1}^{J} p_i^2 ,$$

Where p_c is the portion of units being classified as class c and 1 - p_c is the portion being misclassified.

Because the objective of DTs is to make leaves homogenous, these algorithms are susceptible to overfitting and require regularization. Three common ways to regularize trees are: (1) limit the maximum leaf size, (2) limit the maximum tree depth during training, or (3) prune the tree depth after training.

Decision Trees: Ensemble Methods

A disadvantage of single DTs is that their "greedy" optimization produces unstable solutions. That is, minor modifications to the input data can produce large effects on the tree structure. To combat overfitting, ensemble methods construct many decision trees which then vote on the prediction for each datum. Bootstrap aggregating or "bagging," is an ensemble method that constructs multiple decision trees by repeatedly resampling training data with replacement and generating a consensus prediction. Even with bagging, greedy trees will tend to use the same features for similar decision sequences and thus be correlated. Random forests combine bagging with a feature bagging scheme that forces these greedy trees to explore different decision sequences with other predictive features: at each split, a given tree in the forest can only choose from a random subset of features (Amit and German 1997; Breiman 2001; Ho 1995; Zhu, Zeng, and Kosorock 2015). Because random forests are built from more varied data (bagging) and are forced to explore a broader solution space (feature bagging), the predictions produced by consensus of the forest tend

to have lower variance than a single DT without increasing bias. That is, while predictions for a single tree are sensitive to noise in the training set, the average of many trees is not if the trees are uncorrelated. However, this decrease in variance comes at a cost of interpretability, and ensemble methods are generally viewed as black-box algorithms.

Another popular ensemble method, "boosting," takes a different approach. Boosting algorithms train very deep trees which are low bias but high variance (weak learners), and successively add more weak learners that focus on areas of the data that the current ensemble is doing a poor job of predicting. Gradient boosting machines (GBM) use the gradient of the overall error for the ensemble to decide where to add weak learners (Breiman 1997; Freidman 1999). Bayesian Additive Regression Trees, BART, is a Bayesian "sum-of-trees" model motivated by ensemble methods and boosting algorithms (Chipman 2010). By using priors, this Bayesian approach eliminates many user decisions about regularization, tree depth, and other hyperparameter tuning options.

Unsupervised Learning

Unsupervised learning algorithms identify patterns and associations in the data without help (i.e., labeled examples) to train on. In most cases, these algorithms are used for clustering data into groups, or reducing the dimensionality of data. Unsupervised learning tasks might include identifying different types of individuals on an online dating

site or reducing a dataset of conversational data with 150 possible human gestures down to 10 latent ones. Some common examples of unsupervised learning algorithms include: (a) *k*-means clustering; (b) hierarchical clustering; (c) principal component analysis; (d) t-distributed stochastic neighbor embedding (t-SNE); and (e) unsupervised learning of documents for text analysis.

Unsupervised Learning Algorithms

k-Means Clustering

The *k*-means is among the simplest clustering algorithms. The intuition behind *k*-means is that the data can be clustered by assigning each of n observations into k ($\leq n$) clusters to the nearest centroid (i.e., mean) in the x-dimensional feature space. The *k*-means clustering is then optimized by minimizing the variance within clusters and maximizing the variance between clusters. Researchers initialize a pre-defined number of cluster centroids k in the feature space, then repeat the following two steps until the within-cluster variance is minimized: (1) re-assign each unit to the nearest centroid/mean using some distance metric (e.g. Euclidean distance) and (2) calculate the new means of the units in each cluster as new centroids. A larger k creates smaller groups with more granularity, while a smaller k means larger groups with less granularity, so it may be useful to run the algorithm multiple

times varying this parameter. Silhouette scores are one popular method for evaluating the coherence and separation of clusters. Figure 6 is an example of k-means clustering.

Hierarchical Clustering

Hierarchical clustering generates, as the name indicates, a hierarchy of clusters Agglomerative hierarchical clustering algorithms (e.g. the popular neighbor-joining UPGMA algorithm) take a "bottom-up" approach where each unit begins as an individual and is repeatedly paired with other individual/clusters as researchers move up the hierarchy to a single group. In "top-down" divisive hierarchical clustering, all observations start in one cluster and the splits are made recursively as we move down the hierarchy. Like k-means, researchers must choose a distance metric to decide which clusters should be paired/divided. The Euclidean distance is a common metric for both k-means and hierarchical clustering, but other options include the Manhattan, Mahalanobis, or Hamming distances. Because hierarchical clustering is sensitive to outliers and overfitting, it is common to bootstrap the data to evaluate the stability of clusters.

Network Community Detection

Complex networks are graphical models where nodes are connected to other nodes via edges. These edges represent some relationship between the nodes. If the nodes are

people in a social network, the edges might represent relationships, but other types of relationships can also be represented in a network: scientific papers that are are connected via citations, or documents that are connected by the co-occurrence of words. "Communities" in networks are clusters of nodes that are more densely connected to each other than they are to other nodes, and the discovery of these communities can be thought of as a form of unsupervised clustering. Even for data that cannot conventionally be thought of as a network, community detection can be a useful unsupervised learning approach when there are not many features, but rich relational data. There are many viable approaches to community detection based on modularity-detection, information theory, and label propagation. Some of these algorithms can produce hierarchical communities as well.

Dimensionality Reduction: Principal Components Analysis

Including noisy or highly-correlated features in supervised learning can potentially lead to overfitting. The simplest approach to this problem is to use domain knowledge to remove irrelevant features. Dimensionality is a way to reduce the number of features and focus on the principal features. Dimensionality reduction can be achieved by feature selection and feature extraction. Researchers could perform feature selection by using something like LASSO to identify relevant features and eliminate irrelevant features. Feature extraction creates new features, where the new features are combinations of the old

features. Researchers can use either a linear or non-linear feature extraction procedure. Principal components analysis is an example of linear feature extraction; t-distributed stochastic neighbor embedding is an example of non-linear feature extraction.

Principal component analysis (PCA) extracts new features from data by remapping the data to a smaller space. It is designed to drop the least important features while still retaining the principal ones. PCA's data reduction technique involves creating one or more index variables, the components, from a larger set of observed variables. PCA is used to determine the optimal choice of observed variables for each component and the optimal weights. Typically, this linear transformation is conducted through an eigenvector decomposition of the covariance matrix of the features. In this space, the axes are uncorrelated and each axis explains a percentage of the variance in the dataset. The first principal component explains the greatest proportion of variance, the second explains the second greatest proportion of variance, and so on. Typically, researchers will only use the first few principal components which account for a majority of the variation and discard the rest as noise. PCA is often used in exploratory data analysis by plotting the dataset onto the first two or three principal components to find clusters in the data. Alternatively, the user might use these principal components as latent features in a downstream supervised learning task.

Dimensionality Reduction: t-Distributed Stochastic Neighbor Embedding

Another approach to dimensionality reduction, *t*-distributed stochastic neighbor embedding (t-SNE), enables reduction by non-linear feature extraction. It is a visualization algorithm of clusters for dimensionality reduction, i.e. reducing high-dimensional data in a low-dimensional space such that similar objects are modeled by nearby points. Researchers use t-SNE to map high-dimensional data to a lower-dimensional space and identify patterns in the data by locating observed clusters. The t-SNE algorithm proceeds in two steps. First, it constructs a probability distribution such that similar objects in the high-dimensional space have a high probability of being selected while dissimilar units have a low probability of being selected. Second, it constructs an analogous probability distribution in the low-dimensional space. The similarity metric is commonly the Euclidean distance. This approach, in contrast to PCA, is very computationally intensive. The input features are also no longer identifiable. It is most often used for data exploration and visualization.

Machine Learning Methods and Text Analysis

As the amount of social science data available online grows, text analysis has become an increasingly important part of social science research. The most common uses of machine

learning for text analysis are supervised classification of documents by label, topic, or sentiment, unsupervised clustering of documents to discover topics, name disambiguation, and entity or meaning recognition.

Supervised Labeling of Documents

A common task in text analysis is the automated labeling of documents for topic, sentiment, or class. For example, rather than manually coding 1000 book reviews for political orientation, researchers might want to label just 250 of them and train an algorithm to label the remaining reviews. By using word frequencies as features, many of the supervised learning algorithms described above can be adapted for this task. These algorithms, called "discriminative" models, learn hard decision boundaries between classes given the observed features. In contrast, "generative" models make structural assumptions about the relationships between classes and features by modeling each class as a probability distribution of features. These models are called generative because once trained, the label distribution can be used to generate features consistent with the class probabilistically (i.e., a document of words that co-occur within that class with high probability). Generative models have been particularly successful for classification in small text corpuses because the strong probabilistic assumptions make them robust to outliers in the training data.

The naive Bayes algorithm is one of the simplest supervised generative text classification models. Using Bayes Rule, naïve Bayes models the probability $[P(Y|X)]$ of a document with word distribution X belonging to a class Y as a product of the prior frequency of documents of that class $[P(Y)]$, and the likelihood of the words observed in the document being generated from that class $[P(X|Y)]$. As a concrete example, suppose that the document to be explored is a Yelp review to be classified as positive or negative. The prior probability of the positive class $P(positive)$ would be the frequency of positive reviews in our training set. If the document consisted of only two words "very delicious," the likelihood would be calculated as the probability of those words occurring together in positive sentiment documents. The naïve assumption in naïve Bayes is that researchers can approximate this probability by simply multiplying the probabilities $P(very|positive)$ and $P(delicious|positive)$ together. There are several ways to calculate the probability of words being generated by documents (i.e., the likelihood), although all of them fundamentally assume that a document is an unordered "bag" of words. The simplest way is to compute the probability of a word times being generated from class Y is to compare the number of times that X occurs or does not occur in training documents of class Y (Bernoulli naïve Bayes). An alternative approach instead considers the counts of word in documents of class Y (Categorical/Multinomial naïve Bayes). Once the prior probabilities and

Bernoulli/categorical word distributions for each class have been calculated from the training data, the model can be used to apply labels to a testing dataset.

Unsupervised Topic Modeling

In unsupervised topic modeling, each document is not labeled with a single class, but as a proportional mixture of latent unspecified "topics." For example, a Yelp review might be labeled as 50% Topic 1, 40% Topic 2, and 10% Topic 3. After applying the topic model, the researcher can then go through the corpus and interpret what the topics mean semantically. In this example, Topic 1 might largely correspond with descriptions of the food, Topic 2 might capture sentiment words, and Topic 3 may be about the service. Topic models are useful tools because they allow the researcher to both discover topics or themes in a corpus without a priori assumptions, and to group documents that are topically or thematically similar together. Furthermore, often the only hyperparameter that needs to be chosen is the number of topics.

Topic models differ from supervised learning of documents in several key ways. First, although a document is still represented as a bag of words, not every word in the document must come from the same topic. Second, the topics themselves are categorically distributed. In the popular probabilistic Latent Semantic Analysis algorithm (pLSA), the process can be thought of as follows: For each position in a document, researchers: (1) pick a topic from

the categorical distribution of topics; and (2) generate a word from the categorical distribution of words given that topic. In the Bayesian version of pLSA called Latent Dirichlet Allocation (LDA), researchers have additional priors on the two categorical distributions, improving the performance on small corpuses.

Word Embedding Using Neural Networks

Topic models discover similarities between documents through the global co-occurrence of words between documents. However, the past five years have seen exciting advances in a different approach to finding structure in text data: learning the relationship between words based on their local co-occurrence with other words (i.e., the local context of words no more than 3-10 words away) using a neural network. Word embedding methods train a neural network to correctly predict the 3-5 words surrounding each word over thousands or millions of examples in large corpora. Although this is technically a supervised learning problem, the actual predictions are not actually of interest. What makes these models useful is the function encoded in the hidden layer of this trained network which transforms words into the lower dimensional space. Somewhat surprisingly, vocabulary words represented in this lower-dimensional embedding space (i.e., transformed by the hidden layer of the neural network) are closer to other vocabulary words with similar semantic meaning and/or syntactic usage. This property has been demonstrated by the

ability of word embeddings to solve analogy problems through simple vector arithmetic. For example, subtracting the word-vector for "man" from "king" and adding the word vector for "women" produces a new word vector that is very close to the word vector for "queen." These analogy tests also highlight that the word embedding space contains latent dimensions of meaning where relevant words are aligned along an axis (e.g., man and woman along a gender axis).

Word embeddings in social science have been used for a few purposes. First, they can be used to compare the relative meaning of words to other words across corpuses or across time. For example, researchers could ask how the usage of the word "man" differs between Victorian novels and modern news media. Second, performing dimensionality reduction on the vector space can help identify interesting latent meaning dimensions (e.g. gender) encoded in the corpus. Third, modifications like Doc2Vec of the original word embedding approaches project documents into the same vector space as the vocabulary words. This is not quite the same as discovering topics, but documents with similar language should be closer in vector space, and users can then perform subsequent unsupervised clustering of documents in this space to find groups of documents that are similar. This approach works better than topic modeling for identifying similar documents that are very short in length (e.g. tweets). The two most popular versions of this algorithm are the original method Word2Vec (Mikolov 2013) and GloVe, an approach that combines this local context

approach with global co-occurrence information when producing vectors. However, word embedding methods are very active areas of research, and newer models like BERT use more sophisticated neural networks to capture more nuanced relationships between words than the widely used Word2Vec and GloVe models.

Machine Learning Methods and Causal Inference

Two of the most important developments in social science methodology over the last several decades include advances in causal inference methods and advances in machine learning methods. The intersection of these two paradigms represents an extremely promising direction for empirical social scientific knowledge. Athey (2018) states: "Despite the fascinating examples of 'off-the-shelf' of slightly modified prediction methods, in general ML prediction models are solving fundamentally different problems from much empirical work in social science, which instead focuses on causal inference" (p. 10). A model with high explanatory power that is optimized for prediction might be a relatively poor model, and even inferior to a model with low explanatory power, for causal inference. Social scientists interested in causal questions prefer unbiased estimates of causal effects to accurate prediction of outcomes. Off-the-shelf machine learning algorithms are not designed to attend to key assumptions in a causal model to identify causal effects, such as the key identifying assumption that treatment assignment is unconfounded. Pearl (2018), in the introductory

chapter "Mind over Data," cautions: "While awareness of the need to a causal model has grown leaps and bounds among the sciences, many researchers in artificial intelligence would like to skip the hard step of constructing or acquiring a causal model and rely solely on data for cognitive tasks. The hope … is that the data themselves will guide us to the right answers whenever causal questions come up. I am an outspoken skeptic of this trend because I know how profoundly dumb data are about cause and effects" (p. 16).

Social science researchers, however, are making progress in merging machine learning and causal inference methods. Scholars of causal inference are increasingly adopting approaches from machine learning, while attending to identification strategies from statistics and econometrics, to enhance research in fields such as economics, political science, and sociology. Methods for conditioning on covariates offer a promising intersection between machine learning methods and the goals of estimating causal effects. In order to estimate causal effects without exogenous shocks, we assume selection on observables, or $Y_i(x) \perp W_i = w_i \mid X_i = x_i$ (Imbens and Rubin 2015). Some research has suggested using LASSO for selecting covariates (Belloni, Chernozhukov, and Hansen 2014), and other work has adopted machine learning algorithms for conditioning on covariates. Propensity score weighting and matching on propensity scores has a long history in the causal inference literature. Propensity scores are summary measures of selection into treatment based on observed covariates: $P(x) = P(W_i = w_i \mid X_i = x_i)$. Estimating propensity scores is a prediction problem,

and thus at least on the surface lends itself well to machine learning approaches. Many recent papers consider various strategies for estimating propensity scores based on machine learning algorithms (Lee, Lessler, and Stuart 2010). Classification trees and random forests offers some promising approaches. Iterative procedures that mimic machine learning algorithms but are not black-box are also possible (see Imbens and Rubin (2015) for a description).

The fundamental problem of causal inference is that researchers do not observe potential outcomes in both the treated and control states, and thus there is no ground truth in causal models like there is in predictive models. Adopting machine learning methods for causal inference thus requires changing the objective function. Machine learning methods, however, can be adopted for experimental and observational data settings, instrumental variable models, difference-indifference designs, and regression discontinuity. Estimation of average treatment effects has benefited from a residual balancing approach (Athey et al. 2016) as well as a double machine learning approach (Chernozhukov et al. (2017). Estimation of treatment effect heterogeneity represents an especially promising use of machine learning methods for causal inference. Athey and Imbens (2017) offer a detailed review of a variety of questions that can be addressed relating to treatment effect heterogeneity. One approach developed by Athey and Imbens (2016) is to use decision trees, "causal trees," to partition the feature space according to effect heterogeneity, and estimate

effects within each leaf (see also Brand et al. (2019) for an empirical application using observational data). This approach yields easily interpretable patterns of variation in treatment effects. Wager and Athey (2017) also offer a "causal forest" based on random forests, which is an average of many causal trees where trees differ due to subsampling. Each of these approaches offers new methods by which to uncover effect heterogeneity.

Conclusion

This entry provides an overview of machine learning methods for social science research. We described the origins of computational social science methods. The entry offered as overview of supervised learning methods, including generalized linear models, support vector machines, k-nearest neighbor, artificial neural networks and deep learning, decision trees, and ensemble methods based on decision trees. Several important considerations relevant to supervised learning algorithms, including use of training and test data and cross-validation, loss optimization and evaluation metrics, bias-variance tradeoff, and overfitting and regularization strategies were noted. This was followed by a discussion of unsupervised learning methods, including k-means clustering, hierarchical clustering, principal component analysis, and t-distributed stochastic neighbor embedding. A section on supervised and unsupervised learning for text analysis followed. Finally, a section specifically on machine

learning methods and causal inference concluded. The interface between these two areas is a promising growing area for empirical social science research.

Machine learning methods are not without important limitations for social science research. First, many machine learning algorithms are black-box. Social scientists, however, may prefer white-box models. Relatedly, machine learning methods are not focused on solving estimation problems (Athey 2018). However, in empirical social science research, researchers are often concerned with how effects change when one or more selected covariates are added to or eliminated from the model.

If the data from which a machine learns are biased, so too will be what researchers learn from that data. This is always true in the model estimation practices, but machine learning models make researchers increasingly dependent on the belief that social scientists should be able to effectively predict social behavior. Blau and Duncan (1967) once said "Sociologists are often disappointed in the size of the residual, assuming that this is a measure of their success in 'explaining' the phenomenon under study. They seldom reflect on what it would mean to live in a society where nearly perfect explanation of the dependent variable ... In such a society it would indeed be true that some are 'destined to poverty almost from birth … by the economic status or occupation of their parents.' Others, of course, would be 'destined' to affluence or to modest circumstances. By no effort of their own could they materially alter the course of destiny, nor could any stroke of fortune, good or ill, lead

to an outcome not already in the cards" (p. 174). Researchers should, in other words, not be dismayed that data and models cannot perfectly predict social behavior, nor should researchers want them to. More existential critiques of machine learning methods are related, including the way in which such models effectively amplify researchers' individual affinities and consequently social, economic and cultural inequality and polarization.

With great automation comes great responsibility. Machine learning methods hold tremendous promise and indeed should produce a profound and transformative impact on empirical social science research. Yet this does not eradicate the need to attend to theoretical considerations in our model building strategies.

Further Readings

Athey, Susan and Guido Imbens. (2016). "Recursive portioning for heterogeneous causal effects." *Proceedings of the National Academy of Sciences* 113(27): 7353-7360.

Athey, Susan and Guido Imbens. (2017). "The state of applied econometrics: Causality and policy evaluation." *Journal of Economic Perspectives* 31(2):3-32.

Attewell, Paul, David Monaghan, and Darren Kwong. (2015). *Data mining for the social sciences: An introduction.* University of California Press.

Bishop, Christopher M. (2006). *Pattern recognition and machine learning.* Springer.

Brand, Jennie E., Jiahui Wu, Bernard Koch, and Pablo Geraldo. 2019. "Uncovering Sociological Effect Heterogeneity using Machine-Learning." Working paper.

Breiman, Leo, J. H. Freidman, R. A. Olshen, and C. J. Stone. (1984). *Classification and regression tress.* Monterey: Wadsworth & Brooks Cole.

Domingos, Pedro. (2015). *The master algorithm: How the quest for the ultimate learning machine will remake our world.* Basic Books.

Foster, Ian, Rayid Ghani, Ron S. Jarmin, and Frauke Kreuter, and Julian Lane. (2017). *Big data and social science: A practical guide to methods and tools.* Boca Raton: Taylor & Francis Group.

James, Gareth, Daniela Witten, Trevor Hastie, and Robert Tibshirani. (2017). *An introduction to statistical learning: With applications in R.* Springer.

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. (2016). *Deep learning.* MIT Press. http://www.deeplearningbook.org

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman. (2009). *The elements of statistical learning.* Springer. Available online at http://www-stat.stanford.edu/~tibs/ElemStatLearn/

Kelleher, John D., Brian Mac Namee, and Aoife D'Arcy. (2015). *Fundamentals of machine learning for predictive data analytics: Algorithms, worked examples, and*

case studies. MIT Press.

Murphy, Kevin P. (2012.). *Machine learning: A probabilistic perspective.* MIT Press.

References

Amit, Y. and D. German. (1997). "Shape Quantization and Recognition with Randomized Trees." *Neural Computation* 9(7): 1545-1588.

Belloni, Chernozhukov, and Hansen. (2014). *Journal of Political Economy.*

Blau, Peter M., and Otis Dudley Duncan. 1967. *The American Occupational Structure.* New York: Wiley.

Breiman, Leo. (1997). "Arcing the Edge." *Technical Report 486.* Statistics Department, UC Berkeley.

Breiman, Leo. (2001). "Random Forests." *Machine Learning* 45(1):5-32.

Chipman, Hugh A., Edward I. George, and Robert E. McCulloch. 2010. "BART: Bayesian additive regression trees." *Annals of Applied Statistics* 4(1): 266-298.

Crevier, D. (1993). *AI: The Tumultuous History of the Search for Artificial Intelligence.* Basic Books.

James, G., D. Witten, T. Hastie, and R. Tibshirani. (2013). *An Introduction to Statistical Learning.* New York: Springer.

Freidman, J. H. (1999). "Greedy Function Approximation: A Gradient Boosting Machine." Working paper.

Geman, S., E. Bienenstock, and R. Doursat. (1992). "Neural Networks and the Bias/Variance Dilemma. *Neural Computation* 4*:* 1-58.

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman. (2009). *The Elements of Statistical Learning.* Springer. Available online at http://www-stat.stanford.edu/~tibs/ElemStatLearn/

Ho, Tin Kam. (1995). "Random Decision Forests." *Proceedings of the 3rd International Conference on Document Analysis and Recognition*, Montreal, QC. August 14-16 1995.

Imbens, Guido and Donald Rubin. (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences*. Cambridge: Cambridge University Press.

Lee, Brian K., Justin Lessler, and Elizabeth Stuart. (2010). "Improving propensity score weighing using machine learning." *Statistics in Medicine* 29: 337-346.

Maini, Vishal and Samar Sabri. (2017). *Machine learning for humans.* Available online: https://www.dropbox.com/s/e38nil1dnl7481q/machine_learning.pdf?dl=0

Pearl, Judea. (2018). *The Book of Why*: *The New Science of Cause and Effect.* Basic Books.

Tibshirani, R. 1996. "Regression Shrinkage and Selection via the Lasso." *Journal of the Royal Statistical Society: Series B (Methodological)* 58(1): 267-288.

Wager, Stefan and Susan Athey. 2017. "Estimation and Inference of Heterogenous Treatment Effects using Random Forests." Working paper, July 11.

Zhu, R., D. Zeng, and M. R. Kosorock. (2015). "Reinforcement Learning Trees." *Journal of the American Statistical Association* 110(512): 1770-1784.