



Applied Artificial Intelligence

An International Journal

ISSN: 0883-9514 (Print) 1087-6545 (Online) Journal homepage: <https://www.tandfonline.com/loi/uaai20>

Missing Data Imputation for Supervised Learning

Jason Poulos & Rafael Valle

To cite this article: Jason Poulos & Rafael Valle (2018) Missing Data Imputation for Supervised Learning, Applied Artificial Intelligence, 32:2, 186-196, DOI: [10.1080/08839514.2018.1448143](https://doi.org/10.1080/08839514.2018.1448143)

To link to this article: <https://doi.org/10.1080/08839514.2018.1448143>



Published online: 19 Mar 2018.



Submit your article to this journal [↗](#)



Article views: 2168



View related articles [↗](#)



View Crossmark data [↗](#)



Citing articles: 2 View citing articles [↗](#)



Missing Data Imputation for Supervised Learning

Jason Poulos and Rafael Valle

Departments of Political Science and Electrical Engineering and Computer Sciences, University of California, Berkeley, CA, USA

ABSTRACT

Missing data imputation can help improve the performance of prediction models in situations where missing data hide useful information. This paper compares methods for imputing missing categorical data for supervised classification tasks. We experiment on two machine learning benchmark datasets with missing categorical data, comparing classifiers trained on non-imputed (i.e., one-hot encoded) or imputed data with different levels of additional missing-data perturbation. We show imputation methods can increase predictive accuracy in the presence of missing-data perturbation, which can actually improve prediction accuracy by regularizing the classifier. We achieve results comparable to the state-of-the-art on the Adult dataset with missing-data perturbation and k -nearest-neighbors (k -NN) imputation.


Introduction

Supervised learning has become an increasingly attractive methodology and proven to be effective in social science applications, such as studies of international and civil conflict (Beck, King, and Zeng 2000; Hill and Jones 2014; Muchlinski et al. 2016) and election fraud (Cant and Saiegh 2011; Montgomery et al. 2015). For supervised classification tasks, the objective is to fit a model on labeled training data in order to categorize new examples. However, the ability of researchers to accurately fit a model and yield unbiased estimates may be compromised by missing data.

Our objective is to compare the out-of-sample performance of three popular machine learning classifiers—decision trees, random forests, and artificial neural networks (ANNs)—trained on imputed or non-imputed (i.e., one-hot encoded) machine learning benchmark datasets that contain various degrees of missing-data perturbation. Researchers analyzing survey data typically choose decision trees or random forests for classification tasks,

CONTACT Jason Poulos  poulos@berkeley.edu  Department of Political Science, 210 Barrows Hall #1950, Berkeley, CA 94720-1950, USA.

Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/UAAI.

 Supplemental data for this article can be access on the [publisher's website](#).

largely because these models require neither imputing missing data nor encoding categorical variables, unlike ANNs or other classifiers.

The primary contribution of this paper is to provide guidance to applied researchers on how to handle missing data for supervised learning tasks. First, we show that imputation methods can increase predictive accuracy in the presence of missing-data perturbation. Second, we show that adding missing-data perturbation prior to imputation can actually improve prediction accuracy by regularizing the classifier. We achieve results comparable to the state-of-the-art on the Adult dataset with missing-data perturbation and k -nearest-neighbors (k -NN) imputation. Lastly, we show that classifiers trained on one-hot encoded data generally yield higher predictive accuracy when the data are not additionally perturbed. For example, a simple one-hot encoded random forests outperforms the state-of-the-art on the Congressional Voting Records (CVRs) dataset with no missing-data perturbation.

The following section describes missing data mechanisms and imputation methods; the ‘Experiments’ section describes our experiments on two benchmark datasets and discusses the results; and the ‘Conclusion’ section concludes and offers possibilities for future research.

Missing data and imputation methods

In this section, we describe the missing data mechanisms underlying patterns of missing data common to survey-based datasets. We then review popular methods of handling missing data.

Missing data patterns and mechanisms

It is important to first distinguish between missing data patterns, which describe observed and missing values, and missing data mechanisms, which relate the probability of missingness (Little and Rubin 2014, Chap. 1). Common missing data patterns in surveys typically include unit nonresponse, where a subset of participants do not complete the survey, and item nonresponse, where missing values are concentrated on particular questions. In opinion polls, nonresponse may reflect either refusal to reveal a preference or lack of a preference (De Leeuw, Hox, and Huisman et al. 2003).

Following the notation of Little and Rubin (2014, Chap. 1), let $Y = y_{ij}$ be a $(n \times K)$ dataset with each row $y_i = (y_{i1}, \dots, y_{iK})$ the set of y_{ij} values of feature Y_j for example i . Let Y_{obs} define observed values of Y and Y_{mis} define missing values. Let M define the missing data identity matrix $M = m_{ij}$, where $m_{ij} = 1$ if y_{ij} is missing and $m_{ij} = 0$ if y_{ij} is nonmissing. The missing data

mechanism is missing completely at random (MCAR) if the probability of missingness is independent of the data, or

$$f(M|Y, \phi) = f(M|\phi) \quad \forall Y, \phi,$$

where ϕ denotes unknown parameters. The missing at random (MAR) assumption is less restrictive than MCAR in that the probability of missingness depends only on the observed data, $f(M|Y, \phi) = f(M|Y_{\text{obs}}, \phi)$ for all Y_{mis}, ϕ . The missing not at random (MNAR) assumption is that the probability of missingness may also depend on the unobserved data, $f(M|Y, \phi) = f(M|Y_{\text{mis}}, \phi)$ for all Y_{mis}, ϕ . Researchers typically assume data are MAR, which mitigates the identifiability problems of MNAR because the probability of missingness depends on data that are observed on all individuals (Tsiatis 2007, Chap. 6).

Imputation methods

Complete-case analysis (i.e., discarding examples with missing values) wastes information and biases estimates unless the missing data are MCAR. Since there is no way to distinguish whether the missing data are MCAR or MNAR from the observed data, a natural strategy is to impute missing values and then proceed as if the imputed values are true values. Imputation methods that rely on explicit model assumptions include *mean or mode replacement*, which substitutes missing values with the mean (for quantitative features) or mode (for qualitative features) of the feature vector, and *prediction model* imputation, which replaces missing values with the predicted values from a regression of Y_{mis} on Y_{obs} .

Explicit modeling methods assume the data are MAR while implicit modeling methods, which are algorithmic in nature and rely only on implicit assumptions, generally do not assume the underlying missing data mechanism. Implicit methods include *random replacement*, where an example with missing data is randomly replaced with another complete example randomly sampled, and *hot deck* imputation, where missing values are replaced by “similar” nonmissing values. Hot deck imputation can be implemented by computing the k -NN of an example with missing data and assigning the mode of the k -neighbors to the missing data. Batista and Monard (2003) use this procedure and find k -NN imputation can outperform summary statistic imputation and internal methods used by decision trees to treat missing data.¹

In related work, Silva-Ramrez et al. (2011) compare imputation using ANNs with mean/mode imputation, regression models (logistic regression and multiple linear regression), and hot deck, finding the ANNs model performs the best on datasets with categorical variables.

One-hot encoding

Another natural strategy in dealing with missing data for supervised learning problems is one-hot encoding. Instead of imputing missing data, one-hot encoding creates a binary feature vector that indicates missing values. For categorical features, one-hot encoding simply treats a missing value symbol (e.g., “?”) as a category when the categorical features are binarized. For continuous features, missing values are set to a constant value and a missingness indicator is added to the feature space. One-hot encoding for missing data yields biased estimates when the features are correlated, which is often the case with survey data, even when data are MCAR (Jones 1996).

Experiments

In this section, we describe our experiment on two machine learning benchmark datasets with missing categorical data, comparing three popular classifiers—ANNs, decision trees, and random forests—trained on either one-hot encoded or imputed data with different degrees of MCAR perturbation.

Benchmark datasets

We experiment on two benchmark datasets from the UCI Machine Learning Repository: the Adult dataset and CVRs dataset (Lichman 2013). The Adult dataset contains $N = 48,842$ examples and 14 features (6 continuous and 8 categorical). The prediction task is to determine whether a person makes over \$50,000 a year. The CVRs dataset contains $N = 435$ examples, each the voting record of a member of the U.S. House of Representatives for 16 key roll call votes. The dataset contains 16 categorical features with three possible values: “yea”, “nay”, and missing. The prediction task is to classify party affiliation (Republican or Democrat). In contrast to the Adult dataset, in which only a few features are highly correlated, many of the roll call votes in the CVRs dataset exhibit strong correlations (Figures SM-1 and SM-2).

The state-of-the-art for the Adult dataset is a naive Bayes classifier that achieves a 14.05% generalization error after removing examples with missing values (Kohavi 1996). The CVRs dataset donor claims to achieve a 5–10% error rate using an incremental decision tree algorithm called STAGGER, although it is unknown to the authors what train-test split is used or how missing values are handled (Schlimmer 1987; Schlimmer and Granger 1986).

Patterns of missing data

Uncovering missing data patterns in the datasets will help to identify possible missing data mechanisms and select appropriate imputation methods.

Figure SM-3 analyzes patterns of missing data in the Adult dataset, in which 7% of the examples contain missing values. Missing data in the Adult dataset is due to item nonresponse, as missing values are concentrated in three of the categorical features—*Work class*, *Occupation*, and *Native country*—and no examples contain entirely missing data. It is unlikely that the data are MCAR because observations that are missing in *Work class* are also missing in *Occupation* (about 6% of examples have missing values in both).

Missing values in the CVRs dataset are not simply unknown, but represent values other up-or-down votes, such as voted present, voted present to avoid conflict of interest, and did not vote or otherwise make a position known. Close to half of the CVRs data contains missing values, which are present in every feature (Figure SM-4). About a quarter of missing data is in the feature *South Africa*, which captures a controversial vote to amend the Export Administration Act to bar U.S. exports to South Africa’s apartheid regime. Twelve percent of missing data are in the feature *Water*, which is a water projects authorizations bill, and 7% of missing data rest in the feature *Exports*, which is a tariff bill. The data are unlikely to be MCAR because 12% of the data are missing in just *South Africa* and less than 1% of examples are missing across all features. It is most likely in this case that the CVRs data are MNAR because the probability of missing a vote or voting present on one important bill should not theoretically be influenced by observed votes on other important bills.

Preprocessing

After randomly splitting each dataset (2/3 for training and 1/3 for testing), we perturb the training data so that the proportion of missing values in the set of categorical features Y_{cat} follows $\delta = \{0.1, 0.2, 0.3, 0.4\}$ according to the MCAR mechanism

$$\Pr(M = 1 | Y_{\text{cat}}, \phi) = \delta \text{ for all } Y_{\text{cat}}.$$

We use missing-data perturbation to study the impact of larger amounts of missing data; however, it is also a form of dropout noise that can be used to control overfitting during the training process and improve the generalizability of the model (Wager, Wang, and Liang 2013).

After one-hot encoding the categorical variables in the training data, we implement each of the following imputation techniques, discussed in the “Imputation methods”: k -NN, prediction model (logistic regression, random forests, or SVMs), mode replacement, and random replacement. We then standardize continuous features by subtracting the mean and dividing by the standard deviation of the feature. The test data are preprocessed in the same manner, with the exception that we do not perturb categorical features in the test data.²

Model training

We train three different classifiers on the preprocessed data: decision trees, random forests, and ANNs. The ANNs consist of four layers, each of the two hidden layers having 1024 nodes, and updates with the adaptive learning rate method *Adadelta* (Zeiler 2012). We explore the hyperparameter space—momentum schedule, dropout regularization, and learning rate—using Bayesian optimization (Snoek, Larochelle, and Adams 2012), which selects optimal models using the mean training error rate as our objective function. Figure SM-5 shows the exploration of hyperparameter space during Bayesian optimization for both datasets. Random forests and decision trees are trained with preselected hyperparameters.

Results

We assess the performance of the classifiers in terms of test set error rate on one-hot encoded or imputed data and for various levels of MCAR perturbation. The results on the Adult dataset and CVRs dataset are plotted in Figures 1 and 2, respectively, with error bars representing ± 1 standard deviation from the test error rate.

One-hot-encoded decision trees outperform the state-of-the-art on the CVRs dataset by over 2% (0.027 ± 0.006). The ANNs classifier trained on data imputed with k -NN yields the lowest generalization error (0.144 ± 0.06) on the Adult dataset with 10% of the categorical feature values perturbed, which is comparable to the state-of-the-art. In comparison, a random forests classifier trained on non-perturbed and one-hot encoded data yields a test error rate of 0.152 ± 0.02 . This comparison shows that the classifiers can overfit the data and, in the case of imputed models, perturbation improves prediction accuracy by regularizing the classifier.

Overall, the results show imputation methods can increase predictive accuracy in the presence of missing-data perturbation. For both datasets, one-hot encoded models trained in the absence of perturbation perform as well as imputed models trained on non-perturbed data. In the case of the Adult dataset, imputation clearly improves accuracy in the presence of MCAR-perturbed data. In contrast, each of the three classifiers trained on the one-hot encoded CVRs dataset perform relatively well across different levels of perturbation. The general pattern of results holds when the classifiers are trained on MNAR-perturbed data (Figures SM-6 and SM-7).

Conclusion

This paper compares methods for imputing missing categorical data for supervised classification tasks. We compare the out-of-sample performance

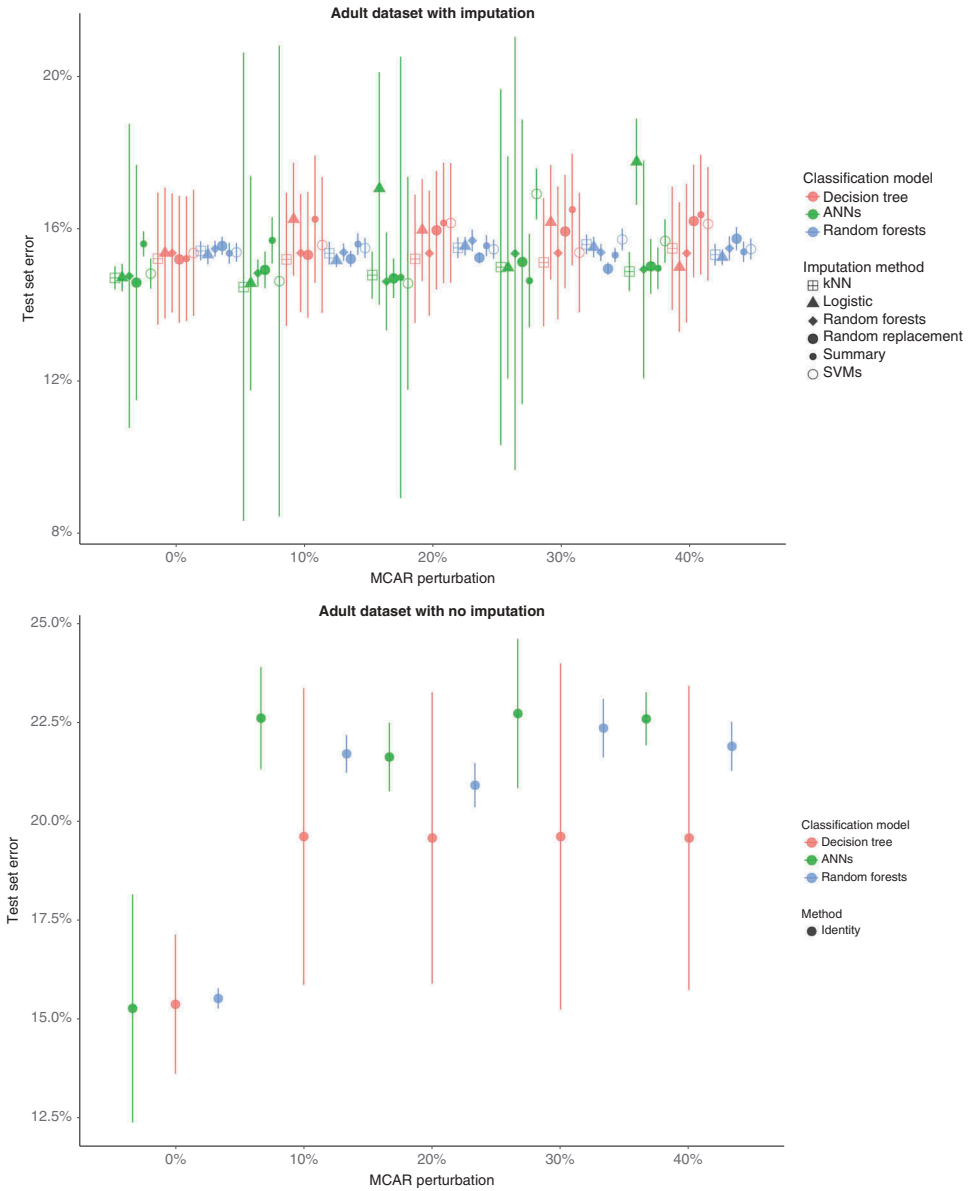


Figure 1. Error rates on the Adult test set with (bottom) and without (top) missing data imputation, for various levels of MCAR-perturbed categorical training features (x -axis). For ANNs, prediction intervals are obtained from the standard deviation of test set errors of ANNs trained with different convergences (Heskes 1997). For random forests and decision trees, prediction intervals follow from the variation created by varying the maximum depth of the decision trees, and for random forests, the number of trees and decision rule for the number of features to consider when looking for the best split.

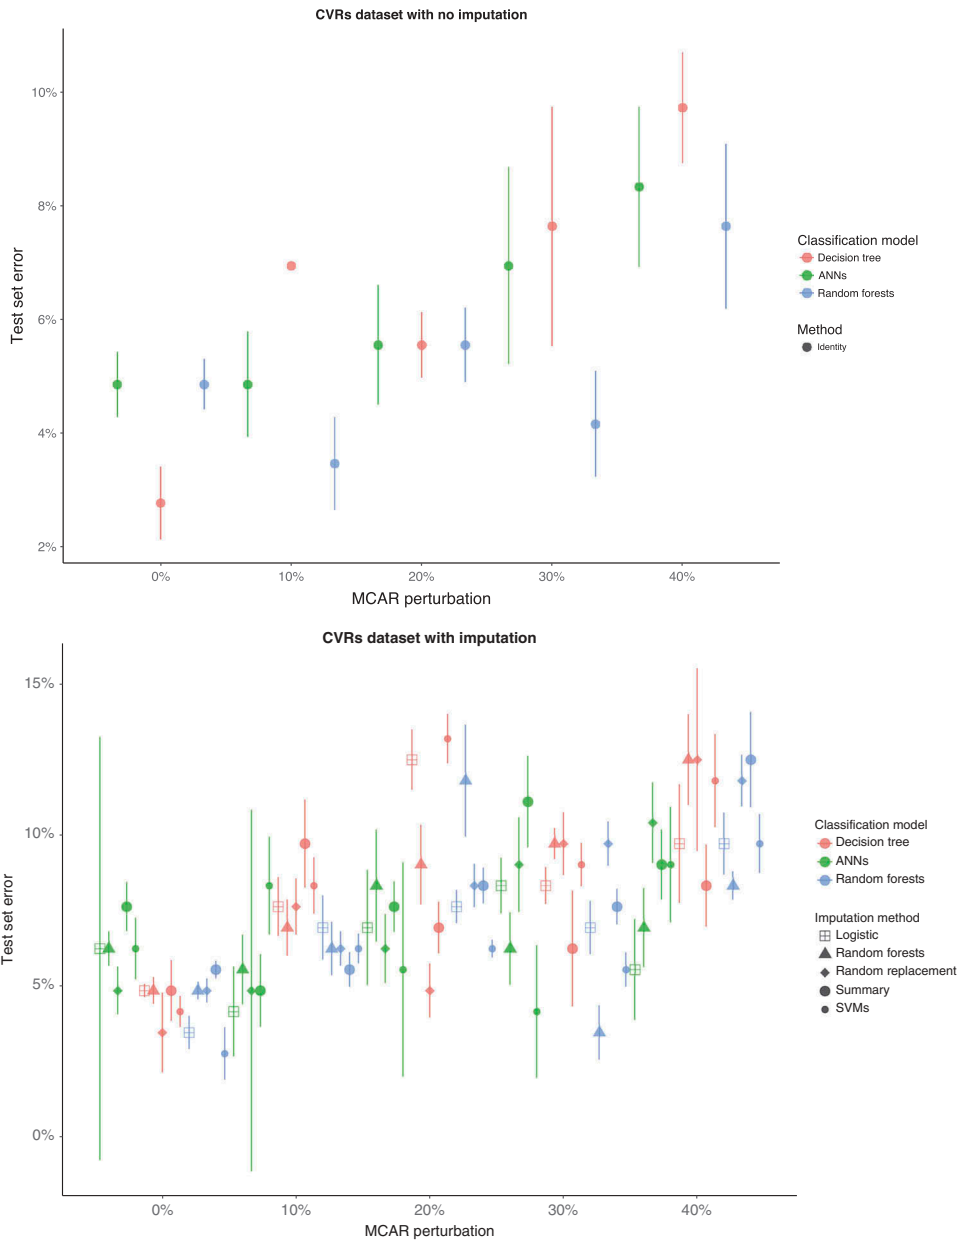


Figure 2. Error rates on the CVRs test set with (bottom) and without (top) missing data imputation. See footnotes for Figure 1.

of ANNs against decision tree and random forest classifiers trained on datasets with one-hot encoded or imputed data, across different levels of MCAR-perturbed training data. Our results are comparable to the state-of-the-art on the Adult dataset using an ANNs classifier, *k*-NN imputation, and

MCAR data-perturbation. k -NN imputation likely performs well in this case because it is an implicit modeling method that does not assume the underlying missing data mechanism. This result is in line with Batista and Monard (2003), who find k -NN imputation can outperform explicit modeling methods for supervised learning tasks.

We conclude from the results that the performance of the classifiers and imputation strategies generally depend on the nature and proportion of missing data. For the Adult dataset, ANNs trained on imputed data generally outperform other classifiers and imputation methods across different ratios of perturbed data, while classifiers trained on one-hot encoded data perform very poorly on perturbed training data.

The results of the present study show that perturbation can help increase predictive accuracy for imputed models, but not one-hot encoded models. Future work can identify the conditions under which missing-data perturbation can improve prediction accuracy. Interesting extensions of this paper include evaluating the benefits of using missing-data perturbation over more popular regularization techniques such as dropout training (Hinton et al. 2012; Maaten et al. 2013; Wang and Manning 2013).

Notes

1. Li et al. (2004) propose a hot deck imputation method based on fuzzy k -means.
2. When imputing the missing data with mode replacement, we use the training set mode. We also use the training set mean and standard deviation to standardize test set features.

Acknowledgments

The code used for this project is available at <https://github.com/rafaelvalle/MDI>. We thank Isabelle Guyon for the idea for the paper and Joan Bruna and seminar participants at the University of California, Berkeley, for comments.

Funding

Poulos acknowledges support from the National Science Foundation Graduate Research Fellowship [grant number DGE 1106400].

References

- Batista, G. E. A. P. A., and M. C. Monard. 2003. An analysis of four missing data treatment methods for supervised learning. *Applied Artificial Intelligence* 17 (5–6):519–33. doi:10.1080/713827181.

- Beck, N., G. King, and L. Zeng. 2000. Improving quantitative studies of international conflict: A conjecture. *American Political Science Review* 94 (1):21–35. doi:10.1017/S0003055400220078.
- Cant, F., and S. M. Saiegh. 2011. Fraudulent democracy? An analysis of Argentina’s infamous decade using supervised machine learning. *Political Analysis* 19 (4):409–33. doi:10.1093/pan/mpr033.
- de Leeuw, E. D., J. Hox, M. Huisman, et al. 2003. Prevention and treatment of item nonresponse. *Journal of Official Statistics* 19 (2):153–76.
- Heskes, T. 1997. Practical confidence and prediction intervals. In *Advances in neural information processing systems*, Vol. 9, ed. M. Mozer, et al., 176–82. Cambridge, MA: MIT Press.
- Hill, D. W., and Z. M. Jones. 2014. An empirical evaluation of explanations for state repression. *American Political Science Review* 108 (3):661–87. doi:10.1017/S0003055414000306.
- Hinton, G. E., N. Srivastava, A. Krizhevsky, I. Sutskever, and R. R. Salakhutdinov. 2012. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*.
- Jones, M. P. 1996. Indicator and stratification methods for missing explanatory variables in multiple linear regression. *Journal of the American Statistical Association* 91 (433):222–30. doi:10.1080/01621459.1996.10476680.
- Kohavi, R. 1996. Scaling up the accuracy of naive-Bayes classifiers: A decision-tree hybrid. In *Proceedings of the second international conference on knowledge discovery and data mining*, ed. E. Simoudis, et al., 202–07. Menlo Park, California: AAAI Press.
- Li, D., J. Deogun, W. Spaulding, and B. Shuart. 2004. Towards missing data imputation: A study of fuzzy k-means clustering method. In *Rough sets and current trends in computing*, ed. S. Tsumoto, et al., 573–79. Berlin, Heidelberg, Germany: Springer.
- Lichman, M. 2013. “UCI Machine Learning Repository.” Accessed October 8, 2015. <http://archive.ics.uci.edu/ml>.
- Little, R. J. A., and D. B. Rubin. 2014. *Statistical Analysis with Missing Data*. Hoboken, NJ: John Wiley & Sons.
- Maaten, L., M. Chen, S. Tyree, and K. Q. Weinberger. 2013. Learning with marginalized corrupted features. In *Proceedings of the 30th international conference on machine learning*, ed. S. Dasgupta and D. McAllester, 410–18. Atlanta, GA: PMLR.
- Montgomery, J. M., S. Olivella, J. D. Potter, and B. F. Crisp. 2015. An informed forensics approach to detecting vote irregularities. *Political Analysis* 23 (4):488–505. doi:10.1093/pan/mpv023.
- Muchlinski, D., D. Siroky, H. Jingrui, and M. Kocher. 2016. Comparing random forest with logistic regression for predicting class-imbalanced civil war onset data. *Political Analysis* 24 (1):87–103. doi:10.1093/pan/mpv024.
- Schlimmer, J. C. 1987. “Concept Acquisition Through Representational Adjustment.” PhD diss., Department of Information and Computer Science, University of California, Irvine.
- Schlimmer, J. C., and R. H. Granger Jr. 1986. Incremental learning from noisy data. *Machine Learning* 1 (3):317–54. doi:10.1007/BF00116895.
- Silva-Ramrez, E.-L., R. Pino-Mejas, M. Lpez-Coello, and C.-D.-L.-V. Mara-Dolores. 2011. Missing value imputation on missing completely at random data using multilayer perceptrons. *Neural Networks* 24 (1):121–29. doi:10.1016/j.neunet.2010.09.008.
- Snoek, J., H. Larochelle, and R. P. Adams. 2012. Practical Bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems* 25, ed. F. Pereira, et al., 2951–59. Red Hook, NY: Curran Associates, Inc.

- Tsiatis, A. 2007. *Semiparametric Theory and Missing Data*. New York, NY: Springer Science & Business Media.
- Wager, S., S. Wang, and P. S. Liang. 2013. Dropout training as adaptive regularization. In *Advances in neural information processing systems*, 26, ed. C. J. C. Burges, et al., 351–59. Red Hook, NY: Curran Associates, Inc.
- Wang, S., and C. Manning. 2013. Fast dropout training. In Proceedings of the 30th international conference on machine learning, ed. S. Dasgupta and D. McAllester, 118–26. Atlanta, GA: PMLR.
- Zeiler, M. D. 2012. ADADELTA: An adaptive learning rate method. *arXiv preprint arXiv:1212.5701*.