PROJECT MUSE®

Considerations Across Three Cultures: Parametric Regressions, Interpretable Algorithms, and Complex Algorithms

Ani Eloyan, Sherri Rose

Observational Studies, Volume 7, Issue 1, 2021, pp. 191-196 (Article)

Published by University of Pennsylvania Press

DOI: <https://doi.org/10.1353/obs.2021.0009>

➔ *For additional information about this article*

<https://muse.jhu.edu/article/799734>

Considerations Across Three Cultures: Parametric Regressions, Interpretable Algorithms, and Complex Algorithms

Ani Eloyan

Department of Biostatistics
Brown University
Providence, RI

ani eloyan@brown.edu

Sherri Rose

*Center for Health Policy
Center for Primary Care and Outcomes
Research Stanford University
Stanford, CA*

sherrirose@stanford.edu

Abstract

We consider an extension of Leo Breiman’s thesis from “Statistical Modeling: The Two Cultures” to include a bifurcation of algorithmic modeling, focusing on parametric regressions, interpretable algorithms, and complex (possibly explainable) algorithms.

Keywords: machine learning, interpretable algorithms, explainable algorithms

1. Introduction

The relevance of the themes presented in “Statistical Modeling: The Two Cultures” by Leo Breiman remains twenty years later. While we could consider many categorizations of statistics culture, we posit that at least three cultures have emerged. We still have the data modeling group with regressions defined within parametric models, but the algorithmic modeling culture has, at a minimum, bifurcated with interpretable algorithms and (possibly explainable) complex algorithms (Rudin, 2019). Practitioners of algorithmic modeling may develop interpretable or complex algorithms or both, depending on the needs of the scientific question. As empirically driven statisticians, we would prefer to collect data on this before putting forth a supposition, but, in lieu of such data, we also surmise that the algorithmic modeling culture has grown larger over time than Breiman’s proposed 2% of statisticians. In this commentary, we remark on several areas of increasing concern for statisticians—in all three cultures—working to solve real, substantive problems.

2. Data

Many statisticians are well aware that an analytic dataset is created after various forms of inclusion and exclusion criteria have been applied in collecting the data, followed by possible dropout of subjects and missingness of variables, among other issues. However, discussions regarding algorithms, such as the choice between parametric statistical approaches and more

flexible machine learning techniques, often omit considerations of more elaborate forms of data preparation, and treat the provided data as observed. Ignoring such preprocessing or data engineering can lead to erroneous results that may be misleading or even harmful (Leek and Peng, 2015; Chen et al., 2021). This is particularly salient with recent forms of data like medical images.

When analyzing functional magnetic resonance imaging (fMRI) data, a collection of preprocessing steps is often applied to prepare data for analyses. Some commonly used preprocessing procedures include motion reduction (as study participants often move in the MRI scanner), transformation of individual images to a common three-dimensional space for group analyses, and image smoothing (Ombao et al., 2016). The choice of preprocessing steps, as well as the selection of software tools for performing preprocessing, can have a tremendous impact on the statistical analyses and resulting inference. For example, using a dataset of images collected from individuals with multiple sclerosis, Eloyan et al. (2014) showed that the results from image registration in a common template space can vary significantly depending on the choice of registration software when individuals are affected by the disease. One possible reason for these differences is that most commonly used software is developed for analyzing data from healthy controls, hence the issues with generalizability to populations with differing brain structures. Using the outputs from these tools as fixed input data as if it were the observed data—for a parametric regression or an algorithm—without deep consideration of the implications of the preprocessing could be disastrous in a health care setting.

3. Algorithms

Analogous to the issues found in overlooking the nuances of data preprocessing, parametric modeling and algorithmic implementations go well beyond selecting an estimator. The reliance on software for applied problems is, again, both a necessary component in the wider adoption of important tools as well as a hindrance to good statistical practice. Various “hidden” assumptions and hyperparameter specifications lurk under the default settings. Here, we examine additional illustrations from imaging.

Statistical parametric mapping is widely used in fMRI studies to understand brain activation, brain functional organization, and differences in brain function (Friston, 1994). These approaches are routinely implemented in software packages used by neuroscientists and clinicians performing the analyses. However, Eklund et al. (2015) demonstrated troubling issues in applying these parametric methods with this off-the-shelf software. Specifically, that statistical parametric mapping methods are invalid for cluster-based fMRI studies. The paper started a wide discussion in the field with calls for change (Brown and Behrmann, 2017; Eklund et al., 2017).

The lack of requisite detail to reproduce analyses is another major problem, particularly in algorithmic implementations. For example, deep learning (Goodfellow et al., 2016) is increasingly used in image analysis for outcome prediction, such as cancer survival or clinical outcomes in dementia. McKinney et al. (2020) proposed a deep learning system making bold claims regarding outperforming radiologists in reading mammograms to identify cancer. Motivated by this work, Haibe-Kains et al. (2020) described issues with transparency and reproducibility in implementations of artificial intelligence systems for imaging-based

prediction studies, e.g., missing hyperparameters from the three algorithms described by [McKinney et al. \(2020\)](#).

These examples show that even if algorithmic approaches have highly accurate performance in a specific study, it may be impossible to replicate the results or implement the same methods in other settings if sufficient details are not provided. Publicly sharing data may not be possible or appropriate due to patient privacy, however, providing similar simulated datasets paired with code can partially resolve this issue, and detailed appendices should include key hyperparameters and other algorithm attributes.

4. Uncertainty and Interpretation

Quantification of uncertainty is often automatic in computational implementations of parametric statistical methods for data analysis, although it is frequently for the wrong target (e.g., standard errors for coefficients rather than prediction intervals for outcomes). Worse yet, in applications of complex or interpretable algorithms, the predicted values or assessment of accuracy is often reported without any estimates of uncertainty of the error. Regardless of which “culture” we fall into, as statisticians, we should be centering the role of uncertainty.

In examples discussed by [Breiman et al. \(2001\)](#), point estimates of prediction or misclassification errors are presented and compared. Bootstrap methods ([Efron, 1992](#)) can be implemented for estimation of prediction and confidence intervals in many nonparametric estimation approaches. However, the implementation of bootstrap methods may not be feasible for some algorithmic approaches due to computational complexity or a breakdown of the required assumptions. While many algorithm-specific methods have been developed for assessing uncertainty in these settings, such as the Bayes-by-backprop approach for neural networks ([Blundell et al., 2015](#)), uncertainty estimates are rarely presented in applications.

Despite the field’s greater focus on algorithms for prediction, machine learning techniques have also been developed for (causal) effect estimation. Uncertainty assessment clearly has major relevance for algorithms designed to estimate effects (what Breiman refers to as “Information”). Here, bootstrapping and influence-curve-based methods have been shown to be useful, such as in targeted maximum likelihood estimation ([van der Laan and Rose, 2011](#); [Cai and van der Laan, 2019](#)). Post-selection inference methods are also available for when model coefficients are the target of inference after a data-adaptive fitting process (e.g., [Berk et al., 2013](#)).

Appropriate interpretations of algorithmic results are crucial. Even with rigorously calculated prediction or confidence intervals, thoughtful consideration of the generalizability or transportability of the study results is essential ([Degtiar and Rose, 2021](#)). To what target population do we aim to make conclusions? Does this tool inform at the population level?

5. Discussion

The statistical literature is often behind other fields involved in data analysis, largely due to the focus on parametric model-based approaches, as discussed by [Breiman et al. \(2001\)](#), including prioritization of “irrelevant theory,” and partially due to the comparatively slower speed of publication in many statistics journals (vs. computer science conference publica-

tions). In addition to expanding our focus on flexible algorithmic approaches, increasing meaningful collaborations with subject matter experts and community members is crucial. Our data are complex, particularly in the area of health (e.g., electronic health records, medical imaging, genomics). It is difficult to learn the necessary background information in isolation in order to form a plausible statistical formulation of the scientific problem and make reasonable assumptions for the analysis methods. We must also consider the potential impact of an algorithm on the communities where it may be applied (Vyas et al., 2020; Chen et al., 2021; Kasy and Abebe, 2021), which also cannot be done in isolation. Statistics as a field has not embedded community-based participatory research in its overall culture to our detriment. What happens next (e.g., after this paper is published in a statistics journal) is typically not asked. Who might use this and how might they use it? Does this tool perpetuate or instigate harm? We argue that these questions, and the other considerations raised in our commentary, are critical areas of focus for all three statistics “cultures.”

Acknowledgments

This work was supported by grant 5P20GM103645 from the National Institute of General Medical Sciences and NIH New Innovator Award DP2MD012722.

References

- Richard Berk, Lawrence Brown, Andreas Buja, Kai Zhang, Linda Zhao, et al. Valid post-selection inference. *The Annals of Statistics*, 41(2):802–837, 2013.
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *International Conference on Machine Learning*, pages 1613–1622. PMLR, 2015.
- Leo Breiman et al. Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical Science*, 16(3):199–231, 2001.
- Emery N Brown and Marlene Behrmann. Controversy in statistical analysis of functional magnetic resonance imaging data. *Proceedings of the National Academy of Sciences*, 114(17):E3368–E3369, 2017.
- Weixin Cai and Mark van der Laan. Nonparametric bootstrap inference for the targeted highly adaptive lasso estimator. *arXiv preprint arXiv:1905.10299*, 2019.
- Irene Chen, Emma Pierson, Sherri Rose, Shalmali Joshi, Kadija Ferryman, and Marzyeh Ghassemi. Ethical machine learning in health care. *Annual Review of Biomedical Data Science*, 2021.
- Irina Degtiar and Sherri Rose. A review of generalizability and transportability. *arXiv preprint arXiv:2102.11904*, 2021.
- Bradley Efron. Bootstrap methods: another look at the jackknife. In *Breakthroughs in statistics*, pages 569–593. Springer, 1992.
- Anders Eklund, Thomas Nichols, and Hans Knutsson. Can parametric statistical methods be trusted for fmri based group studies? *arXiv preprint arXiv:1511.01863*, 2015.
- Anders Eklund, Thomas Nichols, and Hans Knutsson. Reply to brown and behrmann, cox, et al., and kessler et al.: Data and code sharing is the way forward for fmri. *Proceedings of the National Academy of Sciences*, 114(17):E3374–E3375, 2017.
- Ani Eloyan, Haochang Shou, Russell Shinohara, Elizabeth Sweeney, Mary Beth Nebel, Jennifer Cuzzocreo, Peter Calabresi, Daniel Reich, Martin Lindquist, and Ciprian Crainiceanu. Health effects of lesion localization in multiple sclerosis: spatial registration and confounding adjustment. *PloS one*, 9(9):e107263, 2014.
- Karl Friston. Statistical parametric mapping. In *Functional Neuroimaging: Technical Foundations*. Academic Press, 1994.
- Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep Learning*. MIT Press, 2016.
- Benjamin Haibe-Kains, George Alexandru Adam, Ahmed Hosny, Farnoosh Khodakarami, Levi Waldron, Bo Wang, Chris McIntosh, Anna Goldenberg, Anshul Kundaje, Casey Greene, et al. Transparency and reproducibility in artificial intelligence. *Nature*, 586(7829):E14–E16, 2020.

- Maximilian Kasy and Rediet Abebe. Fairness, equality, and power in algorithmic decision-making. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 576–586, 2021.
- Jeffrey Leek and Roger Peng. Statistics: P values are just the tip of the iceberg. *Nature News*, 520(7549):612, 2015.
- Scott Mayer McKinney, Marcin Sieniek, Varun Godbole, Jonathan Godwin, Natasha Antropova, Hutan Ashrafian, Trevor Back, Mary Chesus, Greg Corrado, Ara Darzi, et al. International evaluation of an ai system for breast cancer screening. *Nature*, 577(7788): 89–94, 2020.
- Hernando Ombao, Martin Lindquist, Wesley Thompson, and John Aston. *Handbook of neuroimaging data analysis*. CRC Press, 2016.
- Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019.
- Mark van der Laan and Sherri Rose. *Targeted learning: causal inference for observational and experimental data*. Springer Science & Business Media, 2011.
- Darshali Vyas, Leo Eisenstein, and David Jones. Hidden in plain sight—reconsidering the use of race correction in clinical algorithms, 2020.