PROJECT MUSE®

Breiman's Two Cultures: You Don't Have to Choose Sides

Andrew C. Miller, Nicholas J. Foti, Emily B. Fox

Observational Studies, Volume 7, Issue 1, 2021, pp. 161-169 (Article)

Published by University of Pennsylvania Press

DOI: <https://doi.org/10.1353/obs.2021.0003>

➔ *For additional information about this article*

<https://muse.jhu.edu/article/799728>

Breiman’s two cultures: You don’t have to choose sides

Andrew C. Miller

acmiller@apple.com

Apple
Seattle, WA USA

Nicholas J. Foti

nicholas_foti@apple.com

Apple
Seattle, WA USA

Emily B. Fox

emily_fox@apple.com

Apple
Seattle, WA USA

Abstract

Breiman’s classic paper casts data analysis as a choice between two cultures: *data modelers* and *algorithmic modelers*. Stated broadly, data modelers use simple, interpretable models with well-understood theoretical properties to analyze data. Algorithmic modelers prioritize predictive accuracy and use more flexible function approximations to analyze data. This dichotomy overlooks a third set of models — mechanistic models derived from scientific theories (e.g., ODE/SDE simulators). Mechanistic models encode application-specific scientific knowledge about the data. And while these categories represent extreme points in model space, modern computational and algorithmic tools enable us to interpolate between these points, producing flexible, interpretable, and scientifically-informed hybrids that can enjoy accurate and robust predictions, and resolve issues with data analysis that Breiman describes, such as the Rashomon effect and Occam’s dilemma. Challenges still remain in finding an appropriate point in model space, with many choices on how to compose model components and the degree to which each component informs inferences.

Keywords — *data models, algorithmic models, mechanistic models, hybrid machine learning models, health*

Introduction

In numerous areas, and especially in scientific disciplines, we face the challenge of drawing inferences from limited observations of complex phenomena, complicating the application of standard machine learning tools. For example, health applications present a number of issues that frustrate the formation of robust predictions and inferences. Observations are often limited to a small subset of variables at play in an intricate physiological process. Clinical studies face challenges in recruitment, compliance, and cost, often resulting in small samples. Observational studies can feature heavy confounding, inducing correlations that may disappear in new contexts. Biomarkers can be governed by unknown physiological mechanisms that drive their short- and long-term dynamics. Given such messy data, we are then tasked with the challenge of specifying a model. On the one hand, we want flexible models to capture noisy, complex processes. On the other hand, we need to incorporate an inductive bias to handle limited data or weakly identified scenarios; ideally, this inductive bias leverages our scientific understanding and ensures that predictions are physiologically plausible and clinically sensible. Through this lens of analyzing health data we revisit Breiman’s classic “Two Cultures” paper.

Breiman describes data analysis as a choice between two cultures: *data modelers* and *algorithmic modelers*. Broadly, data modelers use simple and interpretable models with well-understood theoretical properties — e.g., linear Gaussian models. Algorithmic modelers, on the other hand, prioritize predictive accuracy and use more flexible models to find patterns and make predictions — e.g., random forests and neural networks. Despite benefits and shortcomings of both frameworks, Breiman makes a strong case for using algorithmic models to tackle real-world data problems. Algorithmic models are appealing for their flexibility and potential to capture the complex dynamics and relationships at play; however, these models can suffer from under-identifiability, even in the presence of a lot of data, leading to implausible predictions and inferences. Data models, in contrast, can be simpler to fit to data — both computationally and in terms of sample complexity — and provide the allure of interpretability. As Breiman discusses, however, such beneficial properties have limited value if the model is fundamentally ill-suited to describe the data at hand.

Breiman’s dichotomy overlooks an alternative approach to modeling data. In some cases, the underlying data generating process is well-understood and approximately described by a *mechanistic model*. Mechanistic models, derived from scientific theories, encode application-specific knowledge about the causal mechanisms that generate the observed data — e.g., through a set of ordinary or partial differential equations. Where data models embrace simplicity and algorithmic models flexibility, mechanistic models confer an inductive bias informed by domain-specific scientific knowledge. Breiman takes aim at the statistics and machine learning communities, but the inclusion of mechanistic models folds in another community of modelers and data analyzers, namely scientists developing and testing theories.

These three modeling frameworks bring distinct benefits and drawbacks, naturally raising the question, can we enjoy the benefits of each instead of simply choosing one? Growing interest has been in researching and developing such *hybrid models* — models that lie on the continuum between *data*, *algorithmic*, and *mechanistic* models. Due to the tremendous computational and algorithmic advances made in the past 20 years since Breiman’s paper, we are no longer faced with choosing between a few discrete modeling frameworks. Automatic differentiation tools (Abadi et al., 2015; Bradbury et al., 2018; Paszke et al., 2019; Team et al., 2016), approximate inference methods (Blei et al., 2017; Hoffman and Gelman, 2014; Kucukelbir et al., 2016; Ranganath et al., 2014), probabilistic programming languages (Bingham et al., 2019; Carpenter et al., 2017; Salvatier et al., 2016; Tran et al., 2016), and advances in causal inference (Bollen, 2014; Hernán and Robins, 2020; Pearl, 2009) together enable us to specify models drawing from a full menu of tools, including flexible (i.e., algorithmic model), simple (i.e., data model), and highly structured (i.e., mechanistic model) components. This more comprehensive space of model choices grants finer control over issues of interpretability and nuisance variation. By interpolating between the extremes of the cultures, we can more precisely tailor a solution to the particular scientific pursuit at hand.

Hybrid models raise new challenges, however. What variation in our data can be informed by scientific understanding vs. data-driven structure? How do we specify the intersection of components? How can we efficiently perform statistical inference and make predictions? In this note, we examine hybrid models in the statistics and machine learning literature framed within our ontology, highlight the issues that motivate their development, and describe the tools that enable their implementation.

The extreme points of model space

By Breiman’s definition, data models tend to be simple, easy to interpret, and theoretically tractable, but have limited expressivity. This framework describes a wide set of standard statistical tools, including linear regression, logistic regression, autoregressive moving average (ARMA) models (Box et al., 2015), generalized linear models (Nelder and Wedderburn, 1972), and linear mixed models (McCulloch and Neuhaus, 2005). Due to the previously mentioned properties, data models are the go-to methods in many fields, such as econometrics, political science, and other social sciences.

Algorithmic models, by contrast, are highly flexible function approximators that learn predictive features (or representations) from data. Tools in this class include neural networks (Rumelhart et al., 1986), random forests (Breiman, 2001), and more complex deep learning models, long short-term memory networks (LSTMs) (Hochreiter and Schmidhuber, 1997), and transformers (Vaswani et al., 2017). Algorithmic models have undergone a large shift in the last 20 years with deep neural networks becoming a dominant modeling approach due to hardware advances and the availability of large datasets.

Mechanistic models are built from first principles, encoding scientific knowledge about the underlying phenomena that produces the observed data. For example, differential equation models in epidemiology — e.g., the *susceptible-exposed-infectious-removed* (SEIR) framework — describe the dynamics of an infectious disease spreading throughout a population (Anderson et al., 1992; Kermack and McKendrick, 1927). Other examples from the biomedical literature include models of insulin-glucose dynamics in diabetics (Bergman et al., 1981; Dalla Man et al., 2002, 2007, 2014) and cardiovascular function (Kotani et al., 2005; McSharry et al., 2003; Trayanova, 2011). While these models are interpretable to domain experts and encode domain knowledge in fine detail, they can be difficult to fit to data or extend to new types of observations. Additionally, mechanistic models have their own assumptions that must be checked.

These three model classes are not exhaustive. Instead, we consider them a set of building blocks that can be combined in various ways to address a variety of scientific questions and prediction problems. Each with distinct strengths and weaknesses, one model class may not suitably address the complexities of the scientific problem at hand.

Overly simple data models, as Breiman argues, may not have the flexibility to capture all of the subtle variation present in the data required to make good predictions. Algorithmic models can be uninterpretable, making it difficult to explain, audit or critique the predictions from a complex neural network or random forest. Overly flexible models are data-hungry, potentially requiring many more observations than a simpler model to obtain comparable predictions.

One major shortcoming of both data and algorithmic models is that they are, fundamentally, driven by correlations. Predictions are rooted in statistical correlations between inputs and outputs present in the training data. This causal agnosticism, coupled with over-parameterized algorithmic models, leads to what Breiman calls the Rashomon effect: There are many parameter settings that lead to comparable predictive performance. And despite the profound evolution of neural network architectures and learning algorithms, this kind of under-specification is common, even in machine learning systems with access to the largest datasets available (D’Amour et al., 2020). Compounding this issue, we may have access to observational data in one context, but must make predictions in a different, unseen context. In health applications, a common goal is to predict the occurrence of rare events, which, by definition, are not well-represented in the training data. Such extreme events may be out of the observed distribution, so models that rely purely on correlations may struggle to produce accurate predictions in these rare (and potentially critical) situations. Introducing an inductive bias that encodes knowledge of the underlying mechanisms generating the data can constrain inferences and lead to more robust predictions.

Mechanistic models alone, however, can oversimplify complex phenomena, making them overly sensitive to unmodeled nuisance variation. Biomedical models, for instance, may summarize complex metabolic processes with relatively simple dynamics at a coarse level, ignoring sources of uncertainty due to noisy inputs or measurement error. Such models are designed in highly controlled settings, and thus may be unable to cope with variation one might see outside the lab.

Exploring model space — hybrid models

Data analysis requires striking a balance between model flexibility and an appropriate inductive bias, tailored to our goals. Instead of limiting ourselves to one framework, we can explore the space of models between the extremes. Movements in model space can introduce constraining

(a) Conceptual model space. Each extreme represents a preference — simplicity, flexibility, or scientifically informed bias.

(b) T1D hybrid model from Miller et al. (2020). Linear dynamics and a flexible link function model time-varying simulator parameters.

assumptions, or conversely, flexibility. Conceptually depicted in Figure 1a, moving toward data models preferences simplicity; moving toward algorithmic models preferences flexibility; and moving toward mechanistic models preferences a scientifically-informed causal inductive bias.¹ The statistics and machine learning communities have been devising and experimenting with these hybrid models for decades now. And with modern computational tools, interpolation in the space of models between these three extremes has become more common.

Modern machine learning models that blend the data and algorithmic paradigms include the structured variational autoencoder (sVAE) (Johnson et al., 2016), and the output interpretable VAE (oi-VAE) (Ainsworth et al., 2018). These models incorporate highly flexible neural network components together with a “data model” flavor — inducing sparsity or using a structured probabilistic graphical model to describe a learned representation of the data. Arguably, the LASSO (Tibshirani, 1996) lies somewhere between data and algorithmic models, as the procedure entails an algorithmic feature selection step, resulting in a simpler, easier-to-interpret linear model with (in many cases) mature statistical guarantees.

Another path for building more inductive bias into algorithmic models is to explore the algorithmic-mechanistic frontier. One example is the physics-informed neural network, which introduces constraints on activations of the neural network to enforce predictions that are consistent with physical laws Karpadne et al. (2017). Equivariant neural networks are another approach that blend algorithmic and mechanistic models by encoding known symmetries in the data directly into the model (Cohen and Welling, 2016). Convolutional neural networks, arguably, also represent a step from algorithmic to mechanistic models by baking translation invariance into its architecture, compared to less constrained multi-layer perceptron models (Rumelhart et al., 1986). Even Breiman’s description of pruning random tree nodes according to domain knowledge is a step from purely algorithmic models toward a more algorithmic-mechanistic model.

Applications in statistical neuroscience provide examples of models that blend the mechanistic and data paradigms. Time series models of neural image data can be constructed out of data model components (e.g., a linear dynamical system) carefully specified to approximate known biophysical models (e.g., Hodgkin-Huxley models of voltage propagation) to discover functional structure (Johnson et al., 2020; Sun et al., 2019).

1. Disclaimer: Mechanistic models aim to leverage a good approximation to the underlying causal mechanism that generates the data to form predictions. Formal causal inference requires more transparent assumptions to make causal statements, and can use data, algorithmic, and mechanistic modeling tools.

Glucose forecasts and T1D

A recent example from our own research focuses on glucose forecasts for type 1 diabetes (T1D), incorporating input information from a continuous glucose monitor (CGM), diet logs, and insulin delivered (Miller et al., 2020). Our goal was to form physiologically plausible and robust forecasts — that is, sensible predictions in counterfactual contexts (e.g., under different insulin delivery or meal schedules). We found that purely data and algorithmic models, like ARMA time series models and LSTM models, respectively, would often produce physiologically nonsensical counterfactual forecasts — predicting large increases in glucose after a large insulin dose, for example.

Alternatively, mechanistic models that simulate the insulin, glucose, and carbohydrate dynamics in the body have been developed to test insulin dosing strategies (Dalla Man et al., 2002, 2007, 2014). These models have typically been used in simulation studies and developed in highly controlled settings. Crucially, though, these models encode a sophisticated, scientifically-driven inductive bias for how insulin, glucose, and carbohydrates interact in T1D individuals.

However, we found that the T1D simulator alone is insufficient to describe real world CGM data. The simulator, though encoding a physiology-based inductive bias, is too inflexible to handle sources of variation observable in real-world CGM data — person-specific diurnal variation in insulin sensitivity, error in diet logs, and the individual’s movement, among other reasons.

We instead developed a hybrid model — we allowed the simulator parameters to vary in time by introducing a latent linear dynamical system (data model) with a neural network link (algorithmic model) to map the latent variable to physiologically meaningful simulator parameters (mechanistic model). See Figure 1b. Our approach obtained flexibility by allowing simulator parameters to vary in time, yet retained a strong inductive bias by essentially forcing predictions to use diet log and insulin inputs to explain CGM changes. The simplicity of the linear dynamical system prior enabled fine control over the stability of the latent process.

Modeling COVID-19 spread

Another example of hybrid modeling can be found in recent work focused on the COVID-19 pandemic. These models incorporate a new source of potentially relevant data, cell phone-based mobility time-series that describe movement of the population. These hybrid models combine the mechanistic structure of the SEIR framework (Kermack and McKendrick, 1927) with flexible components to incorporate time-varying mobility data to inform the transmission parameters, modulating the basic or effective reproductive number (Arik et al., 2020; Flaxman et al., 2020; Liu et al., 2020; Miller et al., 2020). Some approaches incorporate time-varying indicators that describe certain public policy ordinances, such as school closings or stay-at-home orders, following a data model style (Flaxman et al., 2020). The aim is to accurately predict future outbreaks and understand the effect of certain policies or population-level behavior changes on the spread of the disease.

This hybrid modeling framework is more flexible than the purely mechanistic modeling framework, resulting in better model fit. Even with the incorporation of mechanistic information, care must be taken when using these models to measure causal effects of interventions or produce long-term forecasts. In a thorough simulation study, Soltesz et al. (2020) conclude that the model of Flaxman et al. (2020) is too flexible to reliably measure the impact of multiple non-pharmaceutical interventions in this observational setting, despite using a mechanistic epidemiological model component. Though the mechanistic model component constrains the results of the analysis to a smaller set of hypotheses than a purely algorithmic model, the hybrid model and data may still not pin down a satisfactory answer, requiring either further assumptions (which may or may not be appropriate) or the acquisition of better data.

Conclusion

Developments in computational and algorithmic tools grant us finer control over model details than ever. Expanding Breiman’s ontology to include mechanistic models — and the space of hybrid models — widens the set of questions we can ask and approaches we can take to answer them. These models can incorporate a sophisticated scientific inductive bias (e.g., via a simulator as in the T1D case), but also retain the flexibility of algorithmic models to cope with unforeseen and unmodeled variation, and the simplicity of data models for statistical inference.

For any particular application, a key question is, what is an appropriate operating point in hybrid model space? To answer, we must compose the data, algorithmic, and mechanistic components to achieve a model that addresses the data analysis question at hand. The optimal way to balance these components is an open question, but one that is specific to the goals of the application — e.g., forming robust predictions of rare events. Specific to the mechanistic component there are a number of challenges to address: How do we incorporate domain knowledge that is difficult to simulate, or with unknown structure? Even given a causal model, can the data we observe identify the parameters necessary to form a new prediction or estimate an effect? How do we check or compare mechanistic model components?

Importantly, not all questions can be answered by any given dataset (particularly observational ones). Sometimes, the data simply do not contain the information necessary to answer a particular query. Even if a mechanistic model gives you some traction — i.e., a good set of structural assumptions to build upon — misspecification and unidentifiability can lead you astray. Even within the hybrid modeling framework, bolstered by countless advances in computation, statistics, and machine learning, the simple insight from John Tukey remains timeless: “The data may not contain the answer.” (Tukey, 1986).

References

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>. Software available from tensorflow.org.
- Samuel K Ainsworth, Nicholas J Foti, Adrian KC Lee, and Emily B Fox. oi-VAE: Output interpretable VAEs for nonlinear group factor analysis. In *International Conference on Machine Learning*, pages 119–128. PMLR, 2018.
- Roy M Anderson, B Anderson, and Robert M May. *Infectious Diseases of Humans: Dynamics and Control*. Oxford University Press, 1992.
- Sercan Arik, Chun-Liang Li, Jinsung Yoon, Rajarishi Sinha, Arkady Epshteyn, Long Le, Vikas Menon, Shashank Singh, Leyou Zhang, Martin Nikoltchev, et al. Interpretable sequence learning for COVID-19 forecasting. *Advances in Neural Information Processing Systems*, 33, 2020.
- Richard N Bergman, Lawrence S Phillips, and Claudio Cobelli. Physiologic evaluation of factors controlling glucose tolerance in man: measurement of insulin sensitivity and beta-cell glucose sensitivity from the response to intravenous glucose. *The Journal of Clinical Investigation*, 68(6): 1456–1467, 1981.
- Eli Bingham, Jonathan P. Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul A. Szerlip, Paul Horsfall, and Noah D. Goodman. Pyro: Deep universal probabilistic programming. *Journal of Machine Learning Research*, 20:28:1–28:6, 2019.
- David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- Kenneth A Bollen. *Structural Equations with Latent Variables*, volume 210. John Wiley & Sons, 2014.
- George EP Box, Gwilym M Jenkins, Gregory C Reinsel, and Greta M Ljung. *Time Series Analysis: Forecasting and Control*. John Wiley & Sons, 2015.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Nectra, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018.
- Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.
- Bob Carpenter, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1), 2017.
- Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International Conference on Machine Learning*, pages 2990–2999. PMLR, 2016.
- C Dalla Man, Andrea Caumo, and Claudio Cobelli. The oral glucose minimal model: Estimation of insulin sensitivity from a meal test. *IEEE Transactions on Biomedical Engineering*, 49(5): 419–429, 2002.

- Chiara Dalla Man, Robert A Rizza, and Claudio Cobelli. Meal simulation model of the glucose-insulin system. *IEEE Transactions on Biomedical Engineering*, 54(10):1740–1749, 2007.
- Chiara Dalla Man, Francesco Micheletto, Dayu Lv, Marc Breton, Boris Kovatchev, and Claudio Cobelli. The UVA/PADOVA type 1 diabetes simulator: new features. *Journal of Diabetes Science and Technology*, 8(1):26–34, 2014.
- Alexander D’Amour, Katherine Heller, Dan Moldovan, Ben Adlam, Babak Alipanahi, Alex Beutel, Christina Chen, Jonathan Deaton, Jacob Eisenstein, Matthew D Hoffman, et al. Underspecification presents challenges for credibility in modern machine learning. *arXiv preprint arXiv:2011.03395*, 2020.
- Seth Flaxman, Swapnil Mishra, Axel Gandy, H Juliette T Unwin, Thomas A Mellan, Helen Coupland, Charles Whittaker, Harrison Zhu, Tresnia Berah, Jeffrey W Eaton, et al. Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature*, 584(7820):257–261, 2020.
- Miguel A Hernán and James M Robins. *Causal inference: What if*. CRC Boca Raton, FL, 2020.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- Matthew D Hoffman and Andrew Gelman. The no-u-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(1):1593–1623, 2014.
- Matthew James Johnson, David Duvenaud, Alexander B Wiltschko, Sandeep R Datta, and Ryan P Adams. Composing graphical models with neural networks for structured representations and fast inference. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 2954–2962, 2016.
- Robert Evan Johnson, Scott Linderman, Thomas Panier, Caroline Lei Wee, Erin Song, Kristian Joseph Herrera, Andrew Miller, and Florian Engert. Probabilistic models of larval zebrafish behavior reveal structure on many scales. *Current Biology*, 30(1):70–82, 2020.
- Anuj Karpatne, William Watkins, Jordan Read, and Vipin Kumar. Physics-guided neural networks (PGNN): An application in lake temperature modeling. *arXiv preprint arXiv:1710.11431*, 2017.
- William Ogilvy Kermack and Anderson G McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A*, 115(772):700–721, 1927.
- Kiyoshi Kotani, Zbigniew R Struzik, Kiyoshi Takamasu, H Eugene Stanley, and Yoshiharu Yamamoto. Model for complex heart rate dynamics in health and diseases. *Physical Review E*, 72(4):041904, 2005.
- Alp Kucukelbir, Dustin Tran, Rajesh Ranganath, Andrew Gelman, and David M Blei. Automatic differentiation variational inference. *arXiv preprint arXiv:1603.00788*, 2016.
- Luyang Liu, Sharad Vikram, Junpeng Lao, Xue Ben, Alexander D’Amour, Shawn O’Banion, Mark Sandler, Rif A Saurous, and Matthew D Hoffman. Estimating the changing infection rate of COVID-19 using Bayesian models of mobility. *medRxiv*, 2020.
- Charles E McCulloch and John M Neuhaus. Generalized linear mixed models. *Encyclopedia of Biostatistics*, 4, 2005.
- Patrick E McSharry, Gari D Clifford, Lionel Tarassenko, and Leonard A Smith. A dynamical model for generating synthetic electrocardiogram signals. *IEEE Transactions on Biomedical Engineering*, 50(3):289–294, 2003.

- Andrew C Miller, Nicholas J Foti, and Emily Fox. Learning insulin-glucose dynamics in the wild. In *Machine Learning for Healthcare Conference*, pages 172–197. PMLR, 2020.
- John Ashworth Nelder and Robert WM Wedderburn. Generalized linear models. *Journal of the Royal Statistical Society: Series A (General)*, 135(3):370–384, 1972.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- Judea Pearl. *Causality*. Cambridge University Press, 2009.
- Rajesh Ranganath, Sean Gerrish, and David Blei. Black box variational inference. In *Artificial Intelligence and Statistics*, pages 814–822. PMLR, 2014.
- David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning internal representations by error propagation. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, 1, 1986.
- John Salvatier, Thomas V Wiecki, and Christopher Fonnesbeck. Probabilistic programming in Python using PyMC3. *PeerJ Computer Science*, 2:e55, 2016.
- Kristian Soltesz, Fredrik Gustafsson, Toomas Timpka, Joakim Jaldén, Carl Jidling, Albin Heimer-son, Thomas B Schön, Armin Spreco, Joakim Ekberg, Örjan Dahlström, et al. The effect of interventions on COVID-19. *Nature*, 588(7839):E26–E28, 2020.
- Ruoxi Sun, Scott Linderman, Ian Kinsella, and Liam Paninski. Scalable Bayesian inference of dendritic voltage via spatiotemporal recurrent state space models. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- The Theano Development Team, Rami Al-Rfou, Guillaume Alain, Amjad Almahairi, Christof Angermueller, Dzmitry Bahdanau, Nicolas Ballas, Frédéric Bastien, Justin Bayer, Anatoly Belikov, et al. Theano: A Python framework for fast computation of mathematical expressions. *arXiv preprint arXiv:1605.02688*, 2016.
- Robert Tibshirani. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- Dustin Tran, Alp Kucukelbir, Adji B. Dieng, Maja Rudolph, Dawen Liang, and David M. Blei. Edward: A library for probabilistic modeling, inference, and criticism. *arXiv preprint arXiv:1610.09787*, 2016.
- Natalia A Trayanova. Whole-heart modeling: Applications to cardiac electrophysiology and electromechanics. *Circulation Research*, 108(1):113–128, 2011.
- John W Tukey. Sunset salvo. *The American Statistician*, 40(1):72–76, 1986.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017.