

Sławomir Czech

University of Economics in Katowice

e-mail: slawomir.czech@ue.katowice.pl

ORCID: 0000-0003-3293-6113

EVALUATING THE ROLE OF MACHINE LEARNING IN ECONOMICS: A CUTTING-EDGE ADDITION OR RHETORICAL DEVICE?

Abstract. This paper explores the integration of machine learning into economics and social sciences, assessing its potential impact and limitations. It introduces fundamental machine learning concepts and principles, highlighting the differences between the two disciplines, particularly the focus on causal inference in economics and prediction in machine learning. The paper discusses diverse applications of machine learning, from extracting insights from unstructured data to creating novel indicators and improving predictive accuracy, while also addressing challenges related to data quality, computational efficiency, and data ownership. It emphasizes the importance of standardization, transparency, and ethical considerations in prediction tasks, recognizing that machine learning is a powerful tool but cannot replace economic theory. Ultimately, researchers remain optimistic about the transformative potential of machine learning in reshaping research methodologies and generating new insights in economics and social sciences.

Keywords: machine learning, rhetoric in economics, scientific progress.

Introduction

Machine learning (ML), a prominent subset of artificial intelligence (AI), offers a versatile set of tools that enable machines to learn from experience and adapt autonomously. In the realm of economics, this technology promises to transform the way economists explore complex economic phenomena, predict market trends, and optimize policy decisions. By leveraging vast troves of economic data, machine learning algorithms can unveil previously undetectable patterns, correlations, and nonlinear relationships. Moreover, they can aid in more accurate forecasting, support real-time decision-making, and optimize resource allocation. However, as the allure of machine learning in economics grows, it becomes vital to critically assess its

genuine potential and ascertain whether its integration is a transformative addition to the economist's toolbox or merely a rhetorical device without substantial impact. This essay seeks to address these questions by analyzing the strengths and limitations of machine learning applications in economics and discussing how they can be harnessed effectively to enrich economic research and policy formulation.

The paper is organized as follows. Section two provides an introduction to the fundamental concepts of machine learning, exploring both supervised and unsupervised learning techniques and highlighting key principles essential to understanding ML's functionality. Section three delves into the intersection of machine learning and economics, emphasizing the primary differences between the two paradigms – the focus on understanding and causal inference in economics versus the emphasis on prediction in machine learning. In section four, we explore the diverse applications of machine learning in economics and social sciences, showcasing how ML can yield novel indicators and measurements. Section five sheds light on the benefits and challenges of prediction, delving into controversies surrounding predictive models in economics as well as ethical considerations of leveraging machine learning for decision-making processes. Final section concludes.

Fundamental concepts of machine learning

Machine learning is a subfield of artificial intelligence that focuses on the development of algorithms and statistical models that enable computer systems to learn from datasets and make predictions or decisions based on data without being explicitly programmed for each task (Jiang, 2021; Joshi, 2020). In other words it is the “capability of a machine to imitate intelligent human behavior” (Brown, 2021). Systems based on ML and AI tend to perform complex tasks in a way similar to how people solve problems. Machine learning revolves around the concept of pattern recognition and inference, wherein the machine learns to identify patterns, correlations, and insights from the provided data to improve its performance over time. The core objective of machine learning is to enable machines to adapt and evolve with experience, akin to how humans learn from experience and adjust their behavior accordingly. Thus, as a subset of artificial intelligence, ML is considered as a tool for stimulating human intelligence and decision making. It is also closely related to data science, as it is designed to extract valuable insights and knowledge from data. Data science

provides the framework for data collection, preprocessing, and analysis, while machine learning provides the techniques to build predictive models and derive valuable information from the data. Together, they form a symbiotic relationship that drives the effective utilization of data to solve real-world problems and make informed decisions across various industries and domains.

For a basic comprehension of machine learning algorithms and their corresponding tasks, our attention should be directed towards the fundamental dichotomy existing between supervised and unsupervised learning paradigms (Athey & Imbens, 2019). The former is characterized by the presence of labeled training data. In this learning paradigm, an algorithm is tasked with learning a mapping function that can predict output labels for new, unseen input data. The training data consists of input-output pairs, where the inputs serve as the features, and the corresponding outputs represent the desired labels or target values. During the training phase, the algorithm adjusts its internal parameters based on the provided labeled data, aiming to minimize the discrepancy between its predictions and the true labels. The primary objective of supervised learning is to enable the algorithm to generalize well to unseen examples, thereby making accurate predictions on new, previously unseen data instances.

Unsupervised learning represents a distinct branch of machine learning wherein the algorithm is presented with unlabelled data, devoid of any explicit guidance or target outputs. The primary objective of unsupervised learning is to discover the underlying patterns, structures, or representations within the data without any pre-existing knowledge about the relationships between data points. This type of learning is particularly useful in scenarios where obtaining labeled data may be challenging or expensive. Unsupervised learning algorithms seek to extract meaningful information from the data, such as clustering similar data points together, reducing data dimensionality, or identifying patterns of interest. One of the key challenges in unsupervised learning lies in evaluating the learned representations, as there are no explicit ground truth labels to measure the model's performance against.

Additionally, one should also mention reinforcement learning, which is a paradigm in machine learning that focuses on an agent's interaction with an environment to learn optimal decision-making strategies. Unlike supervised and unsupervised learning, reinforcement learning operates in a dynamic environment where the agent takes actions to maximize cumulative rewards over time. The agent receives feedback in the form of rewards or penalties based on the actions it takes, guiding it towards learning which

actions are favorable in different states of the environment. The primary goal of this approach is to enable the agent to discover an optimal policy that maximizes the expected cumulative reward.

Machine learning encompasses a diverse array of algorithms, each tailored to address specific problems and yield predictions. Among these are regression, support vector machines, decision trees, random forests, K-means clustering, and various other methods. For more informed review of particular methods and approaches refer to the vast literature on ML (e.g. Flach, 2012; Shalev-Shwartz & Ben-David, 2022). In this paper our primary focus lies in elucidating the attainable outcomes and capabilities facilitated by these algorithms. For instance, supervised learning finds common applications in tasks like image classification, natural language processing, and regression, where the model predicts continuous values. On the other hand, unsupervised learning proves valuable in clustering customer segments, detecting anomalies, and performing feature learning for subsequent supervised tasks. By delving into the underlying structure within the data, unsupervised learning algorithms play a crucial role in revealing hidden insights and facilitating a deeper comprehension of intricate datasets.

Several key principles underpin the foundation of machine learning, driving its success and widespread adoption. First and foremost is the principle of generalization, which emphasizes the ability of a trained model to perform well on unseen or new data. A successful machine learning algorithm should not merely memorize the training data but instead grasp underlying patterns to make accurate predictions on novel examples. Secondly, the principle of optimization plays a critical role, as machine learning models seek to minimize or maximize an objective function through iterative optimization algorithms, such as gradient descent. This ensures that the model converges to an optimal set of parameters that yield the best performance on the given task. Additionally, the principle of model evaluation stresses the importance of assessing a model's performance using appropriate metrics to gauge its effectiveness, ensuring its suitability for real-world applications. Moreover, the principle of overfitting avoidance guides researchers to prevent models from becoming overly complex and tailored to the training data, thereby compromising their ability to generalize to new data. Regularization techniques and cross-validation are often employed to mitigate overfitting and enhance generalization capabilities. Overall, these principles form the bedrock of machine learning, enabling the creation of robust, adaptable, and effective models that drive progress across a wide range of industries and scientific domains.

Machine learning and economics

Even though the application of machine learning to economics can be traced back to 1970s (Gogas & Papadimitriou, 2021), it evolved slowly mostly due to technical and data limitations. Low processing powers of early computers made it difficult to process vast amounts of data in reasonable time, not even to mention the lack of systematized and adequate data required by machine learning algorithms. Yet when datasets became increasingly available, the number of connection threads between machine learning and economics flourished and today application of computational approach to economic problems has become one of the most captivating and appealing areas of economics. Undeniably, it was the association between quantitative methodologies in economics and the advancements in machine learning that has blossomed significantly as the potential to integrate the qualitative research with ML was rather limited.

In a broad context, economics primarily concerns itself with comprehending specific social phenomena, emphasizing causal inference to establish connections between causes and effects. Additionally, economics is keenly interested in counterfactual prediction, which involves considering various potential states of reality by altering cases or model parameters. In contrast, machine learning predominantly focuses on predictive tasks, without explicitly addressing causal relationships between variables. Thus it is easy to point differences between economics and computational sciences (Hansen, 2018). Most importantly, classic economic inquiries, including human decision making, market institutions or the nature of company, often pose challenges beyond the capabilities of machine learning techniques to address adequately. For many economists quantitative methods are just tools employed in order to understand the essence of such phenomena. They seek convincing interpretations rather than precise calculations with narrow or restricted meaning (Backhouse, 2011). In effect, machine learning models and algorithms taken off-the-shelf are by and large inappropriate for economists working with this motivation in mind as they typically produce outcomes of a distinct nature when contrasted with the economists' expectations.

Nevertheless, economists devoted to quantitative research have long claimed that the cooperation between econometrics and machine learning has great potential (Varian, 2014). However, they often struggle to cooperate due to fundamental differences in their methodologies and objectives. While econometrics aims for causality and model interpretability with a focus on economic theory, machine learning prioritizes predictive accuracy and com-

plexity, making it challenging to reconcile the two approaches seamlessly. The primary objective of econometrics (and statistics) is inference, aiming to understand how the outcome is connected to the input data. To achieve this, the initial step involves contemplating a model capable of generating such data, followed by the estimation of model parameters using the available data. Although this approach offers relatively straightforward interpretation, it overlooks model uncertainty and the performance of the model on unseen data. Econometricians also focus on estimating specific model parameters to deduce data properties and assess their hypotheses.

In contrast, the core focus of machine learning revolves solely around achieving high prediction accuracy. This entails forecasting outcomes for inputs that were previously unknown during the model's training phase. Emphasizing prediction leads to a preference for complex modeling techniques that excel at handling out-of-sample data. However, this approach often results in black-box outcomes, providing limited insights into how predictions were generated and how they should be interpreted. Unlike traditional concerns for theoretical purity and model elegance, machine learning prioritizes predictive success over theoretical underpinnings. Consequently, the ML workflow involves exploring numerous regressions and models, aiming to identify the most effective one, often with little regard for a carefully chosen theoretical foundation, as commonly observed in econometric practices.

Empirical analysis utilizing the machine learning approach involves estimating and comparing multiple models to identify the one with the optimal prediction accuracy (Athey, 2019). This process often includes tuning parameters and hyperparameters to achieve the best fit for the dataset, although it remains flexible and subject to change when confronted with new sample data. In contrast, the conventional econometric approach centers on methodically selecting a model and estimating its parameters only once. Subsequently, the focus shifts to applying statistical theory to establish confidence intervals for the parameters and possibly exploring counterfactual predictions through parameter manipulation. The primary objective here is to estimate the causal effect of a given phenomenon, guided by a thorough comprehension of the model and a degree of scholarly intuition.

Another distinguishing factor lies in the emphasis on computational efficiency of algorithms. Economists and econometricians typically do not place significant emphasis on this aspect, as their research methods and datasets are often smaller and do not require the scalability demanded by machine learning. ML specialists, on the other hand, prioritize computational efficiency because their algorithms are designed to work swiftly and effectively

with large datasets. Machine learning methods achieve their speed and scalability due to the underlying algorithms that are well-suited for handling vast amounts of data. Economists are content with less efficient algorithms, as their datasets are comparatively smaller and the research methodologies employed differ from those used in machine learning. For example, historical data on past economic crises, which might be of great interest to economists, would be considered negligible for machine learning, which thrives on vast amounts of detailed data to produce accurate predictions. Moreover, many economic phenomena are non-stationary, meaning that predictions based solely on historical data can be biased. Machine learning approaches, in contrast, often assume that the future is a straightforward extension of the past, which might not always hold true in dynamic economic contexts.

A distinct disparity in terminology exists between econometrics and machine learning, even though they may cover very similar concepts. Table 1 highlights how the same categories are labeled differently in the two disciplines. This differentiation in language can serve the purpose of distinguishing the two sciences and fostering easier communication within their respective communities, each with its own specialized terminology. However, this jargon can also act as a tool employed by insiders to create a barrier between themselves and outsiders or “profans.” McCloskey (1983) convincingly argues that the choice of words is a classic rhetorical strategy used to convince others of one’s correctness and to erect a barrier between those who possess knowledge and those who do not. By using specialized terminology and complex language, economists can create a divide that separates their field from the general public and shields them from scrutiny or criticism by those without a deep understanding of economic concepts. This use of jargon can make economic discussions appear more esoteric and exclusive, further reinforcing the perception of economics as a discipline that may seem inaccessible to the broader public. Similarly, in the context of machine learning, the use of technical jargon and intricate language could lead to a comparable effect of creating a non-inviting image.

Table 1
Semantical differences between econometrics and machine learning

	Econometrics	Machine learning
Variable “Y”	Dependent variable, outcome	Label, response, target
Variable “X”	Covariate, explanatory variable, independent variable	Features, predictors
Model building	Estimation	Learning

Source: Hansen (2018, p. 372)

Lastly, one of the challenges with machine learning in economic research is the relationship between data availability and its usability for meaningful analysis. Machine learning methods often rely on existing datasets, which were not purposefully designed by researchers, but were created in order to follow business logic or in line with technological standards. Thus economists should learn how to proactively design data collection processes to align with their research intentions, rather than solely relying on pre-existing datasets that might have a high noise-to-signal ratio. Ensuring data quality and relevance becomes essential for accurate and reliable economic analysis. Furthermore, data collection algorithms and the context in which data is generated (e.g., user experiences that produce customer data) are subject to frequent changes. This dynamic nature of data collection practices raises questions about the robustness of the data and the implications of these changes on the validity of the analysis.

ML applications in social sciences

Despite the aforementioned limitations and reservations, the integration of machine learning methods has sparked a revolution in numerous research domains within the social sciences in recent years, catapulting data-driven analysis to the forefront of scholarly inquiry. The advent of advanced algorithms has empowered researchers to unveil fresh insights into intricate social phenomena, fundamentally reshaping the landscape of empirical investigations. Notably, one of the great advantages of machine learning lies in its capacity to exploit new types of data that were previously obscured and unstructured. Machine learning proves indispensable in making sense of noisy and messy data, facilitating easy information retrieval and pattern discovery. Both supervised and unsupervised methods play vital roles in data exploration. On one hand, they offer richer data descriptions based on numerous dimensions, and on the other, they aid in extracting key features for future utilization. In essence, machine learning possesses the ability to transform unstructured and unexplored data into traceable quantitative forms, rendering them more amenable to traditional statistical analysis.

Thanks to this transformational capability, economists and other social scientists have not only been able to delve deeper into their datasets but also construct new indicators, transcending the limitations of traditional data gathering and interpretation (Mullainathan & Spiess, 2017). For instance, web scraping has enabled the retrieval of millions of retail prices daily

from various global sources, which are then transformed into price indicators for comparison with official Consumer Price Indexes. This method has proven particularly effective in countries where high inflation is concealed through manipulated governmental indicators (Cavallo & Rigobon, 2016). Satellite image processing has revolutionized economic measurement, allowing for the estimation of crop yields, land use, pollution, resource exploitation, and areas of increased or diminished economic activity (Donaldson & Storeygard, 2016). It has also opened new possibilities for tracking poverty in both developing and developed countries using satellite images, Google Street View, and data from mobile phone networks (Blumenstock, 2016). Machine learning algorithms have been successfully applied in financial market analysis by deriving sentiment and uncertainty from data collected from tweets and leading financial newspapers (Baker et al., 2016). In historical research, machine learning algorithms have been harnessed to measure individual mobility during significant events such as the Great Depression (Feigenbaum, 2015).

In certain instances, economic decision-making can effectively leverage the machine learning emphasis on prediction accuracy. Chalfin et al. (2016) presented a compelling argument that the integration of machine learning tools may contribute to increased marginal labor productivity. McBride and Nichols (2018) demonstrated that these methods can lead to improved targeting of poverty policies. Consequently, economists find it more conducive to identifying optimal policies or theories that better reflect reality. It is important to note that econometrics remains crucial in this context, as it provides the necessary statistical and causal inference framework. This incorporation of machine learning not only introduces new data and methodologies but also prompts the exploration of novel questions and avenues of inquiry within economics. An extensive body of research based on anomaly detection approaches has been effectively harnessed by banks to combat credit card frauds (Bin Sulaiman et al., 2022). The adoption of machine learning techniques has thus proven instrumental in enhancing decision-making processes, policy targeting, and problem-solving in various economic domains.

Scholars from other social sciences have also recognized the promising potential of machine learning as a valuable research tool. For instance, sociologists have adopted algorithmic techniques to analyze media discourses, state documents, community creation, and network activities (Molina & Garip, 2019). This approach enables them to uncover population heterogeneity by identifying clusters based on hidden structures and symbolic boundaries, which may have gone unnoticed through traditional

statistical analysis. Machine learning proves especially valuable in capturing nuanced patterns within complex social phenomena. In the legal domain, it has proven beneficial in the analysis of contracts and legal documents, as well as in law practice and administration (Surden, 2021). Not only does it aid in sifting through thousands of documents to extract relevant information, but it also plays a role in estimating the probability of winning a case or understanding the implications of contract terms. Moreover, government agencies, such as courts and police, leverage machine learning for bail decisions, sentencing considerations, and the assessment of the likelihood of reoffending. These applications enhance decision-making processes and contribute to the efficiency and accuracy of legal proceedings.

Yet there are compelling arguments put forth by scholars indicating that machine learning will never fully replace human understanding of complex social problems. The benefits of human interpretation, description, and contextual knowledge are essential in grasping the nuances and complexities inherent in social phenomena. The challenges of data and case selection, as well as the inherent difficulty in establishing causal inference, pose significant obstacles for purely algorithmic approaches (Grimmer, 2015; Hindman, 2015). From this perspective, social scientists are likely to encounter many of the long-standing concerns that have troubled quantitative scholars, such as data quality and preprocessing, computational complexity, sampling bias, and reproducibility issues. The incorporation of machine learning in the social sciences does not eliminate the need for rigorous data handling and validation, nor does it bypass the complexities of identifying causality in social systems. However, this convergence of challenges and concerns opens up a promising opportunity for more efficient collaboration between scholars from various disciplinary backgrounds and areas of expertise. By combining the strengths of human insight, domain knowledge, and interpretative abilities with the computational power of machine learning, researchers can forge a synergistic approach that tackles complex problems in novel and insightful ways. This integration of diverse perspectives can lead to more robust research outcomes, facilitate interdisciplinary cooperation, and enhance the overall advancement of knowledge in the social sciences.

Benefits and challenges of prediction

The inherent nature of machine learning methods, as discussed earlier, has sparked enthusiasm among economists to employ them in prediction tasks. In many policy applications Kleinberg et al. (2015) argue that causal

inference is not always the central problem. Often, the prediction of policy consequences outweighs the need to uncover true causes of specific phenomena, such as targeting social interventions or assessing the probability of committing a crime. Successful prediction can, in fact, provide valuable insights into underlying causes and shed light on theoretical issues more effectively than traditional empirical studies.

The growing computational nature of social science is beginning to shift the traditional bias against prediction, as highlighted by Hoffman et al. (2017). The common objection that predictive accuracy leads to complex, uninterpretable models that lack generalizability and insight is being reevaluated. The authors propose that predictive accuracy and interpretability need not be mutually exclusive but can be complementary. Making predictive models more interpretable can involve using simpler models, focusing on essential predictors, and utilizing visualizations to aid researchers and stakeholders in understanding the model's predictions. Conversely, making interpretive models more predictive can entail using more sophisticated statistical techniques, integrating diverse data sources, and employing machine learning algorithms to uncover complex patterns in the data.

However, there are crucial issues that need resolution before machine learning-based prediction can become more useful and scholarly in the social sciences with higher replicability (Athey, 2019). Standardization of prediction evaluation needs improvement, with a preference for studies that use agreed-upon standardized benchmarks or preregister their predictions in confirmatory studies. Recognizing and popularizing the theoretical limits of predictive accuracy in complex social systems is essential, distinguishing highly regular events from wildly unpredictable ones (the so-called black swans). Before embracing prediction and interpretation as complementary, a better understanding of their respective limitations is crucial. Only through careful consideration and addressing these challenges can machine learning-based prediction be leveraged effectively in the social sciences, leading to more robust and insightful research outcomes.

Athey (2017) highlights critical considerations that arise from the direct implementation of statistical models in practice. As decision-making becomes increasingly reliant on computers and algorithms, such as in self-driving cars or medical treatment strategies, the emergence of "machine ethics" becomes inevitable (Moor, 2006). Stakeholders may not always comprehend the reasons behind specific algorithmic decisions, posing significant challenges in resource allocation and medical commitments. Poor transparency and limited interpretability of algorithm results can lead to skepti-

cism and apprehension about the widespread use of machine learning techniques. Additionally, the issue of fairness and discrimination in artificial intelligence presents a complex problem that cannot be simply reduced to numerical metrics. Human-related values play a vital role in determining what is considered fair or discriminatory and can vary depending on context, timing, and location (John-Mathews et al., 2022).

Finally, the ownership of vast datasets predominantly lies with private companies, which may be reluctant to share them with scholars and the open public. These companies often prioritize commercial exploitation of their data, which may hinder its availability for research and public decision-making purposes. This lack of access to critical data can hinder the advancement of knowledge and potentially limit the public's ability to benefit from the potential insights that could be derived from such data.

Conclusion

For some economists, especially those with qualitative leanings, machine learning may be perceived as a novel element of economic rhetoric that does not address fundamental questions in economics. However, for quantitative economists, it represents a revolutionary paradigm. Despite grappling with challenges, they are determined to overcome them. It needs to be highlighted though, that while machine learning offers a potent and computationally predictive tool, it remains just that – a tool. It cannot replace the indispensability of economic theory. Machine learning's potential lies in its ability to aid in data-driven model selection, as highlighted by Athey (2019, p. 535). She foresees a profound transformation of economics through artificial intelligence and ML, with new research programs, questions, and improved measurement (p. 542). It is undeniable that many empirical-oriented researchers will embrace machine learning methods, driven by sound reasons. The scope of machine learning in public administration, particularly in healthcare and education, is expected to broaden significantly. However, this expansion will necessitate the establishment of new standards for research documentation and experiment replicability, which will likely be overcome in due time.

Machine learning has the potential to automate many routine and mundane data analysis tasks, streamlining the research process and allowing economists to focus on more intricate aspects of their work. It can revolutionize decision-making processes, foster interdisciplinary collaborations, and offer novel insights into complex social and economic phenomena. While

it cannot replace economic theory, it is destined to play a significant role in reshaping research methodologies, generating new questions, and facilitating better measurement in economics and beyond. As researchers continue to confront challenges, they remain optimistic that the benefits of machine learning will outweigh the obstacles, leading to a new era of data-driven innovation and understanding in the field of economics.

REFERENCES

- Athey, S. (2017). Beyond prediction: Using big data for policy problems. *Science*, 355(6324), 483–485. <https://doi.org/10.1126/science.aal4321>
- Athey, S. (2019). The Impact of Machine Learning on Economics. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *Chicago scholarship online. The economics of artificial intelligence: An agenda* (pp. 507–547). The University of Chicago Press.
- Athey, S., & Imbens, G. W. (2019). Machine Learning Methods That Economists Should Know About. *Annual Review of Economics*, 11(1), 685–725. <https://doi.org/10.1146/annurev-economics-080217-053433>
- Backhouse, R. E. (2011). *The puzzle of modern economics: Science or ideology*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511780196>
- Baker, S. R., Bloom, N., & Davis, S. J. (2016). Measuring Economic Policy Uncertainty *The Quarterly Journal of Economics*, 131(4), 1593–1636. <https://doi.org/10.1093/qje/qjw024>
- Bin Sulaiman, R., Schetinin, V., & Sant, P. (2022). Review of Machine Learning Approach on Credit Card Fraud Detection. *Human-Centric Intelligent Systems*, 2(1–2), 55–68. <https://doi.org/10.1007/s44230-022-00004-0>
- Blumenstock, J. E. (2016). Fighting poverty with data. *Science*, 353(6301), 753–754. <https://doi.org/10.1126/science.aah5217>
- Brown, S. (2021). *Machine learning, explained*. MIT Management Sloan School. <https://mitsloan.mit.edu/ideas-made-to-matter/machine-learning-explained>. Retrieved 17.07.2023.
- Cavallo, A., & Rigobon, R. (2016). The Billion Prices Project: Using Online Prices for Measurement and Research. *Journal of Economic Perspectives*, 30(2), 151–178. <https://doi.org/10.1257/jep.30.2.151>
- Chalfin, A., Danieli, O., Hillis, A., Jelveh, Z., Luca, M., Ludwig, J., & Mul-lainathan, S. (2016). Productivity and Selection of Human Capital with Machine Learning. *The American Economic Review*, 106(5), 124–127. <https://doi.org/10.1257/aer.p20161029>
- Donaldson, D., & Storeygard, A. (2016). The View from Above: Applications of Satellite Data in Economics. *Journal of Economic Perspectives*, 30(4), 171–198. <https://doi.org/10.1257/jep.30.4.171>

- Feigenbaum, J. J. (2015). *Intergenerational Mobility during the Great Depression*. Harvard University Working Paper.
- Flach, P. A. (2012). *Machine learning: The art and science of algorithms that make sense of data*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511973000>
- Gogas, P., & Papadimitriou, T. (2021). Machine Learning in Economics and Finance. *Computational Economics*, 57(1), 1–4. <https://doi.org/10.1007/s10614-021-10094-w>
- Grimmer, J. (2015). We Are All Social Scientists Now: How Big Data, Machine Learning, and Causal Inference Work Together. *Political Science & Politics*, 48(01), 80–83. <https://doi.org/10.1017/s1049096514001784>
- Hansen, S. (2018). Machine Learning for Economics and Policy. In J. J. Ganuza & G. Llobet (Eds.), *Social and Economics Studies: Vol. 5. Economic Analysis of the Digital Revolution* (pp. 369–397). Funcas.
- Hindman, M. (2015). Building Better Models. *The ANNALS of the American Academy of Political and Social Science*, 659(1), 48–62. <https://doi.org/10.1177/0002716215570279>
- Hofman, J. M., Sharma, A., & Watts, D. J. (2017). Prediction and explanation in social systems. *Science*, 355(6324), 486–488. <https://doi.org/10.1126/science.aal3856>
- Jiang, H. (2021). *Machine learning fundamentals: A concise introduction*. Cambridge University Press. <https://doi.org/10.1017/9781108938051>
- John-Mathews, J.-M., Cardon, D., & Balagué, C. (2022). From Reality to World. A Critical Perspective on AI Fairness. *Journal of Business Ethics*, 178(4), 945–959. <https://doi.org/10.1007/s10551-022-05055-8>
- Joshi, A. V. (2020). *Machine learning and artificial intelligence*. Springer Nature. <https://doi.org/10.1007/978-3-030-26622-6>
- Kleinberg, J., Ludwig, J., Mullainathan, S., & Obermeyer, Z. (2015). Prediction Policy Problems. *The American Economic Review*, 105(5), 491–495. <https://doi.org/10.1257/aer.p20151023>
- McBride, L., & Nichols, A. (2018). Retooling Poverty Targeting Using Out-of-Sample Validation and Machine Learning. *The World Bank Economic Review*, 32(3), 531–550. <https://doi.org/10.1093/wber/lhw056>
- McCloskey, D. N. (1983). The Rhetoric of Economics. *Journal of Economic Literature*, 21(2), 481–517. <http://www.jstor.org/stable/2724987>
- Molina, M., & Garip, F. (2019). Machine Learning for Sociology. *Annual Review of Sociology*, 45(1), 27–45. <https://doi.org/10.1146/annurev-soc-073117-041106>
- Moor, J. H. (2006). The Nature, Importance, and Difficulty of Machine Ethics. *IEEE Intelligent Systems*, 21(4), 18–21. <https://doi.org/10.1109/mis.2006.80>
- Mullainathan, S., & Spiess, J. (2017). Machine Learning: An Applied Econometric Approach. *Journal of Economic Perspectives*, 31(2), 87–106. <https://doi.org/10.1257/jep.31.2.87>

- Shalev-Shwartz, S., & Ben-David, S. (2022). *Understanding machine learning: From theory to algorithms*. Cambridge University Press. <https://doi.org/10.1017/CBO9781107298019>
- Surden, H. (2021). Machine learning and law: An overview. In R. Vogl (Ed.), *Research handbooks in information law. Research handbook on big data law* (pp. 171–184). Edward Elgar Publishing Limited.
- Varian, H. R. (2014). Big Data: New Tricks for Econometrics. *Journal of Economic Perspectives*, 28(2), 3–28. <https://doi.org/10.1257/jep.28.2.3>