

TECHNOLOGY

Improving Performance in the World of AI: The Rise of Prediction Products

by Pietro Micheli and Galit Shmueli



Image Credit | Aan

AI prediction products: improving predictions to fit reality or modifying reality to fulfill prediction?

✔ **INSIGHT** | FRONTIER 11 Feb 2025

In increasingly turbulent and dynamic business environments, managers are seeking tools that provide reliable data, generate accurate predictions, and support them in making better decisions. Existing approaches draw on traditional performance measurement and management tools (KPIs, performance targets, scorecards) and on big data analytics (Sanders, 2016). However, recent uses of Artificial Intelligence (AI), especially by online platforms, are creating novel scenarios in relation to the acquisition, analysis and use of data with the aim of predicting human behavior and organizational performance with potentially radically new consequences for individuals and organizations.

RELATED CMR ARTICLES

Ångström, R. C., Björn, M., Dahlander, L., Mähring, M., & Wallin, M. W. (2023), “**Getting AI Implementation Right: Insights from a Global Survey**,” California Management Review, 66/1: 5-22.

Kolbjørnsrud, V. (2024), “**Designing the Intelligent Organization: Six Principles for Human-AI Collaboration**,” California Management Review, 66/2: 44-64.

In this frontier article, we focus on a specific use of AI tools – the creation of “prediction products” – that has recently arisen thanks to platformization, whereby platforms facilitate, aggregate, and monetize interactions between end-users and providers of products and services. So far, researchers and managers have mainly concentrated on how platforms use AI to understand and predict human behavior via machine learning algorithms that interact with users. For example, at Netflix machine learning algorithms are deployed to understand viewers’ preferences and to provide recommendations to engage them with the movies they most prefer. Here, we consider situations where three parties are involved: a platform (first party) collects data on users (second party) and generates and sells predictions to a buyer (third party), typically an organization. In this context, the platform monetizes users’ data not by directly sharing it with the buyer, but

rather by “packaging” the data, using AI, into a “prediction product” (Zuboff, 2019), that is, the prediction of users’ behaviors. In the next sections, we start by briefly reviewing the more established approaches and then concentrate on the AI-based one, highlighting various crucial issues. We conclude by discussing the main implications for managers.

Performance Measurement and Management: Measure-to-Modify

Traditionally, organizations have adopted a hypothesis-driven approach to the process of data collection, analysis, and communication where purposes such as monitoring, improving and predicting performance, are defined up front. In practice, strategic objectives are articulated and several tools, such as scorecards, KPIs and targets, are introduced to both control and direct efforts towards the attainment of the strategic objectives. To further motivate employees to achieve them, a system of incentives is often implemented in the form of financial and non-financial rewards (Micheli and Manzoni, 2010).

This approach is mainly top-down: strategic objectives are identified first and then mechanisms are put in place to promote their achievement; sometimes the objectives are changed or refined because of previous results. Importantly, even though the collection, analysis and communication of data are often regarded as “neutral”, this process is “performative”: what is being measured changes because of the very fact of being measured. Hence, we label this approach “measure-to-modify”. For example, if a KPI is introduced to evaluate employee productivity, such as to measure how many patients are seen by a doctor at a hospital, it is likely that employees will try to improve their productivity, whether positively – by shortening patients’ waiting times and improving their pathways – or negatively – by seeing more patients but eventually achieving worse clinical outcomes, because assessments are too brief.

When deploying performance measurement tools, organizations have increasingly attempted to predict, rather than simply monitor, their performance, for example by introducing leading, rather than lagging indicators. In this context, prediction refers to using a KPI to increase the chances of achieving a certain goal, as the KPI could provide

insight and greater confidence that a goal would be attained. This confidence is based on the hypothesized causality between attributes and depends on an organization's ability to create and refine causal models. For example, in a production plant, if setup times are reduced then it can be reasonably hypothesized that lead times will reduce too. In this paradigm, predictive ability could be also enhanced by combining measurement with other tools such as scenario planning and forecasts.

Big Data Analytics: Predict-to-Modify

Thanks to the increasing availability of data from a variety of sources, a second approach is becoming commonplace. In this case, while performance measurement tools may still be used, the emphasis is more on an organization's capacity to leverage data gathered from different sources (e.g., internal data, sensors, wearable devices) to predict future performance using algorithms (Sanders, 2016). For example, we have worked with companies like GE, Siemens, and Rolls-Royce that have invested in predictive maintenance whereby data on the wear and tear of their products is gathered to predict breakdowns or malfunctions. Similarly, in many companies the allocation of employees across stores or work sites now depends on the demand or workload predicted using past data. Data can also be gathered to predict customer choices; for example, based on which genres and songs a person has listened to, Spotify offers suggestions through personalized playlists. These are all examples of predictive analytics which consist of examining historical data, detecting relationships, and extrapolating to formulate predictions.

We label this approach "predict-to-modify" because organizations formulate predictions with the aim of modifying their work processes or offerings - in the examples above, their maintenance schedules, employees' workloads, and users' playlists. While sharing some similarities with the "measure-to-modify" approach, in these contexts, data may be gathered in a variety of ways and perhaps with different purposes but are then used to identify patterns and correlations, and to formulate predictions. In this case, there may be no targets, incentives or intentional feedback loops: the performative element is less pronounced, and individuals may be aware that data are being collected, but not why.

At the same time, tools and practices from the two approaches can be used jointly: various authors have emphasized that the combination of big data analytics and performance measurement tools such as KPIs can be a key way to provide prediction, “not just rearview-mirror reviews” (Schrage and Kiron, 2018). For example, customer loyalty may be hypothesized as a primary driver of sales volumes; however, data may show that investments to increase customer loyalty do not significantly impact sales. Therefore, a more inductive approach might be adopted to identify drivers of sales, for example by conducting small experiments in different markets. This could then lead to a reconceptualization of the causal model – which addresses the question, “what drives sales?” – that could then result into further hypothesis testing.

AI-based Approach: Predict-then-Modify

In this frontier article, drawing on secondary data and on experience we gained by researching digital platforms, such as Google, Amazon and Meta, we concentrate on a third type of approach, which makes it possible for platforms to generate almost perfect predictions (“prediction products”) for their business customers (Shmueli and Tafti, 2023). For example, Google’s “predictive audiences” tool aimed at e-commerce customers uses machine learning to “predict the future behavior of your users”¹ and provides business customers with lists of users who are most likely to purchase, churn, or generate the most revenue in the near future,² such as over the next 7, 14 or 28 days. More generally, in our research we have come across a wide range of predicted user behaviors which can be used for a variety of purposes, including insurance companies using predicted user behaviors to tailor personalized insurance products, advertisers using predicted user preferences to offer personalized ads, and political campaign managers predicting user voting intentions to target suitable audiences (Shmueli and Tafti, 2023).³

When offering prediction products, the seller aims to showcase high prediction accuracy and the buyer expects that this is achieved via smart algorithms applied to quality data. While prediction accuracy can indeed be improved by using larger, more robust datasets and by enhancing the quality of the algorithms, it can also be achieved by combining AI tools with “persuasive technology”, which refers to behavior modification techniques implemented on online platforms (Greene *et al.*, 2022). In this context, behavior may refer

to clicking on an ad, purchasing an item, and posting information. Behavior modification techniques, such as “nudging”, refer to ways of manipulating choice environments to subtly guide users towards making certain choices. For instance, nudges are built into the design of recommender systems on e-commerce platforms such as Amazon, so that “personalized product recommendations” and “relevant search” results nudge users to purchase more and/or specific products in order to generate profit.

While nudging and similar techniques have been used in business for years, the use of persuasive technology on online platforms is more recent and often implemented via various machine learning algorithms that operate in a data-driven, autonomous, interactive, and sequentially adaptive manner. This means that the combination of prediction and behavior modification – that we label “predict-then-modify” – may result in a black box system where the goal is to have perfect alignment between predicted and real values. We describe the system as a “black box”, because how this is attained is not clear to platform users, customers, or even the platform data scientists. In the words of former Facebook Platform Operations Manager, Sandy Parakilas, “these systems have been built in such a way that they’re hard to control and optimize. I would argue that we humans are now out of control. We’ve built a system that we don’t fully understand.”⁴

What is particularly problematic is that it is in the platforms’ interest to push users’ behaviors towards the algorithmic predictions they generate to showcase its predictive accuracy and ability to influence behaviors. We regard this as problematic, because pushing outcomes towards predictions can be harmful to manipulated users and at odds even with companies’ intentions and goals. For instance, a platform like Amazon could sell predictions of customer purchases to companies offering products on the platform (e.g., “Amazon sellers”). These companies could use the prediction products they purchased to adjust their production, marketing and sales efforts and strategies. However, a dysfunctional effect would be that, for customers predicted to be unlikely to purchase these companies’ products, the platform could achieve high prediction accuracy by actively dissuading these customers from purchasing the product, by promoting the products less to these users or by recommending to them other products or sellers, even if these are less relevant for them. In all these cases, it is evident how the “predict-then-modify” approach can have negative consequences not only for the second party (users), but also for the third (buyers).

Comparing this approach to the ones outlined earlier, we can see some similarities, but also substantial differences (see Table 1). For example, in both the “measure-to-modify” and the AI-driven “predict-then-modify” cases there is a performative element—we want to modify behaviors and therefore feedback loops are present—there are three key differences. First, in the AI scenario individuals are not aware that data is being collected, by whom or why. Second, in the “predict-then-modify” case, the data is typically owned by the seller (platform) whereas in the “measure-to-modify” one the organization owns the data. Third, in the world of AI, unexpected consequences are easily exacerbated because of faster feedback loops.

Moreover, in the “predict-then-modify” approach, the end goal of achieving similarity between prediction and reality is not in the usual sense of improving predictions to fit reality; rather, vice versa: modifying reality to fulfill prediction. In this setting, what is being created is neither a model nor a “neutral prediction” anymore, as the algorithm’s main role is not to represent reality, but to modify it. Predictions therefore serve as self-fulfilling prophecies. Unlike the first scenario, where performance indicators are intended to be known, the algorithmic predictions are not revealed to any of the parties but are used by the platforms’ automated persuasive technology to modify behaviors toward those predictions. This means that the performative element is not clearly owned, and responsibility and accountability are opaque. Indeed, we could say that the algorithm is in control, but the algorithm itself is a black box.

	Performance measurement and management “Measure-to-modify”	Big data analytics “Predict-to-modify”	AI-driven “Predict-then-modify”
How predictions are generated	By using tools such as leading KPIs, identified by hypothesizing cause-and-effect relationships, and combined with forecasts and scenarios	By deploying algorithms to examine historical data, detect relationships, and extrapolate to formulate predictions	By deploying automated black box machine learning algorithms
Causality	Causal models are created up front and refined through data analysis over time	Both correlation-based (using interpretable predictive analytics) and causal-oriented (using experiments with manually selected treatments)	Correlation-based predictions, then automated causal behavioral modification algorithms
Awareness of behavior modification by subjects, feedback loops	Performance indicators are typically made known to encourage behavioral effects and feedback is purposefully provided (performativity)	The performative element is less pronounced and there may be no intentional feedback loops; individuals may be aware that data are being collected, but not why	AI overtly uses behavior modification based on fast feedback loops (persuasive technology) without users’ awareness or consent
Behavioral effects	High, mainly predictable	Limited	High: by changing users’ preferences, they become more predictable; however, they are predictable by the algorithm but invisible to users, organizations (buyers of prediction products), and platform managers
Potential unintended effects	Gaming, tunnel vision	Limited	Pushing platform users toward harmful behaviors, without their knowledge or consent; misleading buyers of predictive products

Responsibility and accountability / Locus of control	Generally clear: from senior management to employees	Employees deploying algorithms, e.g., data science teams	The algorithm (black box)
------------------------------------------------------	------------------------------------------------------	----------------------------------------------------------	---------------------------

Table 1: Comparing different approaches to collecting, analyzing and using data to modify behaviors and generate predictions

Conclusions and Implications for Managers

In recent years, we have witnessed a substantial change in the uses of performance data to generate predictions in organizational contexts. In this article, we have discussed the main differences between three approaches and their varied effects on the relationship between prediction and behavior modification. The case of “prediction products” highlights a new capability that has recently emerged due to platformization: predicting and modifying human behavior via algorithms that interact with users.

This predict-then-modify scenario has significant managerial implications. First of all, since the quality of the service delivered by platforms is mainly expressed as the similarity between predicted and actual behaviors, it is imperative that managers understand how platforms attain such similarity and do not simply trust the basic similarity scores. While identifying prediction products that use persuasive technology to improve their accuracy is not easy, managers could ask providers for greater transparency as to the algorithms used, and explicitly probe whether and how behaviors are being manipulated. For example, if one of the aims is to maximize user engagement, how is the prediction being generated?

Furthermore, as data ownership and accountability are shifting from being business-owned to platform-owned in various industries, the AI-driven approach discussed here creates new relationships between firms and platforms. Platforms use a variety of data monetization strategies that managers should be aware of. Such strategies may offer new opportunities to firms but can also create unintended harms. For example, Uber now monetizes its hugely rich data on drivers, riders, and routes to offer business customers targeted advertising to riders. Fitbit, the provider of wearable technology, physical fitness monitors and activity trackers, uses its users’ data to offer a Health and Wellness Insights

Platform that provides anonymized, aggregated data to healthcare providers, insurers, and researchers. eBay monetizes data by offering retailers, manufacturers, and industry analysts access to aggregated and anonymized sales data.⁵ In this article, we have discussed the known examples of prediction products sold by Google and Facebook. Uber, Fitbit, eBay and others could potentially create prediction products as well and sell them to business customers. For example, Uber might sell predicted risk scores of drivers to insurance companies. Fitbit could sell predicted health risk scores of its wearable users to healthcare providers, and pharmaceutical and insurance companies. eBay could sell predictions of its sellers' behaviors, including chances of successful future transactions to third parties such as financial services firms. Managers should assess the potential consequences of platforms creating and selling these products. Also, they should consider up front whether the prediction they are asking for is related to any specific risk. For instance, is there a risk of overly engaging users (especially if young people)? Could addictive or divisive behaviors be promoted? What is the likelihood of putting users or society at risk? Reflecting on these aspects at the very beginning would help managers frame the focus of the prediction in such a way that, however generated, it would be less likely to lead to unintended consequences at a later stage.

References

Greene, T., Martens, D. and Shmueli, G. (2022). Barriers to academic data science research in the new realm of algorithmic behaviour modification by digital platforms. *Nature Machine Intelligence*, 4(4), 323-330.

Micheli, P. and Manzoni, J.-F. (2010). Strategic performance measurement: Benefits, limitations and paradoxes. *Long Range Planning*, 43, 465-476.

Sanders, N. (2016). How to use big data to drive your supply chain. *California Management Review*, 58(3), 26-48.

Schrage, M. and Kiron, D. (2018). Leading with next-generation key performance indicators. *MIT Sloan Management Review*, Research report:

<https://sloanreview.mit.edu/projects/leading-with-next-generation-key-performance-indicators/>

Shmueli, G. and Tafti, A. (2023). How to “improve” prediction using behavior modification”. *International Journal of Forecasting*, 39(2), 541-555.

Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*, Profile Books.

1. <https://support.google.com/analytics/answer/9846734>
 2. <https://support.google.com/analytics/answer/9805833>
 3. <https://theintercept.com/2018/04/13/facebook-advertising-data-artificial-intelligence-ai>
 4. <https://www.pbs.org/wgbh/frontline/interview/sandy-parakilas/>
 5. <https://www.datacommercecloud.com/post/data-monetization-examples>
-



Pietro Micheli [Follow](#)

Pietro Micheli is Professor of Business Performance and Innovation at Warwick Business School. His research and consulting interests are at the intersection of strategy, innovation, and operations management. He has worked with over 60 firms and published in leading journals, including the Academy of Management Journal, Research Policy and others.



Galit Shmueli [Follow](#)

Galit Shmueli, Chair Professor at National Tsing Hua University, researches statistical and machine learning methods in information systems and healthcare, focusing on prediction and causation. Author of popular textbooks and 100+ publications, her work appears in leading journals like Management Science, Marketing Science, Information Systems Research, and Statistical Science.