

Who Do We Blame?

Ayan Bhattacharya

Baruch College, The City University of New York
New York, USA.

© 2018 Ayan Bhattacharya, All Rights Reserved.

For the past few weeks, India has been fixated on the Nirav Modi-PNB saga—with every passing day exposing some new or unexpected affront of the regulatory system. Investigative authorities have finally begun swooping in on the details, and like every time, there is the hope that at least this time, the guilty would be brought to book. At the heart of this, or, for that matter, any other investigation lies a basic assumption: that the sleuths, judges, juries, authorities—and even suspect—understand and agree on the rules of the game. But imagine a world where no one fully grasps the rules underpinning the system—this, increasingly, is the world of financial markets manned by sophisticated, intelligent algorithms. Now if something goes wrong in such a world, who do we blame?

1. Learning the Rules of the Game

How do societal rules arise? This question has a long, rich history that has spawned entire fields of study, like Sociology. Within economics too, there is a vibrant tradition of research on this question. While different schools of inquiry differ on the finer details, at the broad level, there is consensus that the rules are a form of context-dependent equilibrium that help us to do well as a community.[1] Think of it this way: if we did not have traffic lights, there would be many more accidents. Traffic lights also work because we drive cars on flat 2-dimensional surfaces— that is the context. If and when we all start driving drones in 3-dimensions, our traffic system will have to evolve. Because social rules are conditional on context, they depend crucially on our learning and state of knowledge. If we did not have the technology for cars, and our only mode of transport was walking, we would not need traffic lights! The world that we inhabit has not changed much in the past few hundred years, yet our social rules have evolved enormously because we have learned more and more about ourselves and the world.

Learning is so innate in us humans that we seldom marvel at the enormous complexity of the process. Researchers, for many decades, have struggled to understand and replicate the intricate cadence that underlies human learning; yet even now, our grasp of the process is shaky. In the last few years, however, we seem to have uncovered a few basic principles that underpin practical learning. Whether these are the same principles that animate human learning, no one knows for sure; yet there is increasing evidence that these principles do produce a flavor of learning. In particular, two important ideas lie behind the explosion in algorithmic learning applications in the last few years: *deep-learning*, and, increasingly, *self-learning*.

2. Learning Deeply, On Your Own

A popular technique to increase one's skill at chess, in the early stages, is to play against oneself. The idea is that as your own opponent, you will create some small variation, to which you will learn to respond optimally – increasing your feel for the game. It is this intuitive idea that motivates the principle of *self-learning*. If the space of strategies of each player is sufficiently well-behaved mathematically, it is possible in-principle to learn the game fairly well by playing against oneself.[2] For example, think of a simple “Guess an integer” game, where you and your friend sequentially announce a positive integer between one and one billion; each integer that is announced has to be higher than all previously announced integer; and the winner is the one who can make the last announcement. You don't really need to play against a friend to understand how best to play this game. Very mechanically, you can assume simple variations for your friend's play (say, for instance, you can assume that your friend will always guess one higher than you—so when you guess 2, he will guess 3) and exhaust all possibilities for the game. The big problem, however, with such a brute force mechanical approach to self-learning is that it takes an awful lot of time. Even for mathematically well-behaved strategy spaces, the time bottleneck is impractical. For self-learning to work, therefore, we need insights about the data that can cut through the fluff and guide us on how we select our hypothetical opponent's strategy.

The constellation of techniques that go under the name of *deep-learning* has been around since the 1960s, but crucial breakthroughs beginning the late-2000s made the approach practical (for example, [3]). The basic idea is to have multiple layers of connected nodes (like human neurons) in a hierarchy, with each layer focused on a certain level of abstraction, and the output of a lower layer serving as input for the subsequent higher layer. So given a game of chess, the lowest layer might recognize just the board and pieces, the next layer might use this input to identify legal moves in the game, the layer after that might begin to recognize “good” moves, and so on. As each further layer in the hierarchy recognizes more and more abstract features, the strength of inter-connections in the lower layers are re-adjusted to better reflect the overall interpretation. Gradually, as the process gorges on more and more data—adding and subtracting inter-connections in various layers on the way—the network begins to “understand” the game. At the big picture level, deep learning is a technique to generate deep insights that requires copious amounts of data.

Now let's put the two and two together. Self-learning generates a lot of data but needs insights about the data to work successfully. Deep learning generates insights but needs bountiful quantities of data to work well. So what is the obvious conclusion? Self-learning and deep learning seem just made for each other! Well, not really, because the real world has innumerable stochastic variables that are highly unpredictable.[4] Think, for example, of a sudden pothole that may emerge in the path of a self-driving car due to unanticipated rain the previous evening. Nevertheless, in certain special settings, the match between self-learning and

deep-learning is indeed strong. These are settings where the rules of engagement are clearly defined, and the majority of players use similar algorithmic techniques. In other words, financial markets!

3. The Financial Market Black Box

If there is one field outside of computer science where algorithmic techniques are having an outsize impact, it is finance. Financial markets present a relatively manageable, controlled environment, which is nevertheless sufficiently rich to present many interesting challenges. Not surprisingly, outside of Silicon Valley, Wall Street firms are among the biggest recruiters of tech talent.[5] At the same time, recent techniques in artificial intelligence and algorithms build on the tools of game theory. Game theory, incidentally, has been the bread and butter of serious economists for many decades. Thus a two-way street seems to have opened up between computer science and economics & finance that is gradually changing the contours of both fields.

Going back to our topic at hand, in a market where traders are self-taught, deep-learning networks, if something goes wrong, how do we know who to blame? In order to assign guilt, we need to understand the motives of the guilty. But self-learning/deep-learning networks are closed loop black boxes where we neither understand how the data for learning is generated, nor how the insights about trading decisions are arrived at. Self-learning and deep-learning, in tandem, are almost a self-regulating structure that brooks no outside intervention. It is quite possible that market outcomes which appear deviant to us, might in fact be stepping stones to smart trades. But we have no way to know. Going back to our old example, algorithms might be capable of 3-dimensional drone maneuvers while we are still stuck in our primitive 2-dimensional traffic lights.

The rules that we have defined for our markets reflect our human capacity for learning. If such markets were inhabited by artificially intelligent algorithms (as they increasingly are), how must we create the rules – especially when algorithmic learning is a black box to us?

4. The Opportunity

Strangely, though the present algorithmic setting is completely novel, this is not the first time humans have grappled with such questions. Ancient traders faced the same question when they landed on an unfamiliar shore, and you and I face the same question when we adopt our first pet. In fact, whenever two distinct cultures of learning come into contact for the first time, we almost always grapple with such questions. Though it appears far removed, the clues to a solution to our ongoing algorithmic conundrum might lie in such encounters. From history, we have learned that exchanges between adherents of distinct styles of learning have proved most successful when there has been no imposition. Instead, what works is a shared system of ethics,

and a commonly accepted system of values within which everyone operates. Thus we have concepts like democracy, privacy, morals and universal rights. A nascent movement in the algorithmic community towards these ideas is already underway.[6] Going back to our traffic analogy, what we really care about is no accidents, not the specifics of any particular traffic system.

Most of these ideas are still in their infancy and only vaguely understood at the present time. As we map out this new and unfamiliar algorithmic terrain, a lot of academic and real world fortunes will be made. After all, every one of us wants a Nirav Modi to stand trial, even if he were an algorithm!

--

[1] Herbert Gintis, "The Bounds of Reason. Game Theory and the Unification of the Behavioral Sciences," Princeton University Press, 2009.

[2] Noam Brown, Tuomas Sandholm, "Superhuman AI for heads-up no-limit poker: Libratus beats top professionals," Science, January 26, 2018.

[3] Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. NiPS12, pp. 1097-1105.

[4] Joshua Sokol, "Why Self-Taught Artificial Intelligence Has Trouble With the Real World", Quanta Magazine, February 21, 2018.

[5] Nanette Byrnes, "As Goldman Embraces Automation, Even the Masters of the Universe Are Threatened," MIT Technology Review, February 07, 2017.

[6] Kevin Hartnett, "How to Force Our Machines to Play Fair", Quanta Magazine, November 23, 2016.