UC Law Business Journal

Volume 20 | Number 2 Article 5

7-2024

The Ethics Of Artificial Intelligence

Justice Gordon Goodman

THE ETHICS OF ARTIFICIAL INTELLIGENCE

*Justice Gordon Goodman**

TABLE OF CONTENTS

INTRODUCTION

The question of how to adopt and use artificial intelligence in ways that are helpful to mankind is now being explored within every industry and by all branches of government. The following article recounts an early attempt to adopt and use a primitive form of artificial intelligence in the energy industry, which at the time was typically called "machine-aided learning." Many of the same problems we encountered, along with similar concerns and forms of resistance, appear in current media reports. I hope these experiences from over twenty years ago will prove instructive to today's "early" adopters of artificial intelligence.

THE SANTA FE INSTITUTE AND COMPLEXITY SCIENCE

During the late 1990s, I served as E.I. DuPont de Nemours & Co.'s representative on the Santa Fe Institute's (SFI) business network. SFI was founded in 1984 to study the emerging field of complexity, and it was led

* Justice Gordon Goodman was elected to the First Texas Court of Appeals in 2018. He is a member of the Texas State Bar, Pennsylvania Bar, and Energy Bar Associations. Before serving on the bench, he started his career with the Whittenburg Law Firm in Amarillo, TX, where he focused on civil trials, appellate work, oil & gas law, banking law, and general practice. He subsequently served in senior positions for the Howell Corporation; E.I. DuPont de Nemours & Co.; Conoco, Inc.; Occidental Petroleum Corporation; and most recently with NRG Energy. He earned his BA degree Magna Cum Laude from the University of Pennsylvania and his JD degree from the University of Pennsylvania Law School, both in Philadelphia. He received his high school degree from the Horace Mann School in the Bronx, NY.

during its early years by Dr. Murray Gell-Mann[1], a Nobel-prize physicist; Dr. Brian Arthur[2], an economist who wrote seminal works on the impact of positive feedback or increasing returns in economies; Dr. John Holland[3], a computer scientist who developed genetic algorithms and learning classifier systems; and other senior scholars. It arose in part from the non-linear dynamics studied at the nearby Los Alamos National Laboratory during development of the hydrogen bomb in the 1950s. SFI's research spans many historically separate disciplines, but it always applies rigorous logic and mathematics to all its studies.

SFI's business network is composed of companies that apply insights from complex systems to their work in many fields (e.g., finance, energy, banking, manufacturing, etc.). While participating in the network's activities, I invited scientists from SFI to visit and lecture at DuPont's Experimental Station in Wilmington, Delaware.

When I joined Occidental Petroleum Corporation (Occidental) in the spring of 1999, I continued my association with SFI. Working closely with Occidental's talented quantitative risk analysts, financial experts, and energy marketers, I began developing a complex, adaptive, predictive-pricing program based on many of the ideas I first encountered at SFI. Within Occidental, we called this program the "Black Box," and it was initially written during the period April-July 1999 at Occidental.

The development of the Black Box program began with our review of standard energy marketing efforts at that time, which were heavily focused on short-term, prompt month, price fluctuations. An overview of the market indicated that only a small percentage of all the profits that could be made from energy marketing were associated with this short-term, prompt month activity. However, it accounted for most of the time and effort by marketing groups within Occidental (and this was not unique to Occidental). It was critically important for Occidental to maximize the price it received from the sale of its energy products and to make wise hedging decisions to avoid future downside price risk.

Most profits from energy marketing were associated with a relatively small number of significant price changes that occurred infrequently. To the extent that we could reliably identify precursory indicators for these larger, long-term, price changes, we could improve results from energy marketing (while engaging in a smaller number of transactions), maximize prices received from crude oil and natural gas sales, and mitigate future price risk through hedging. The question was whether it was possible and how to identify these precursory events in the marketplace using existing data.

1. Dr. Murray Gell-Mann authored a well-known popular work on complexity titled *The Quark and the Jaguar: Adventures in the Simple and the Complex* (Freeman & Co., New York, 1994).

2. Dr. Brian Arthur wrote *Increasing Returns and Path Dependence in the Economy* (University of Michigan Press, Ann Arbor, 1994).

3. Dr. John Holland authored the foundational work on genetic algorithms titled *Adaptation in Natural and Artificial Systems* (MIT Press, 1975). He also authored a more popular discussion of these ideas titled *Hidden Order: How Adaptation Builds Complexity* (Addison-Wesley, New York, NY., 1995)

DEVELOPMENT OF THE BLACK BOX PREDICTIVE PRICING PROGRAM

The following are the basic assumptions that we applied in writing the initial Black Box predictive-pricing program, which was preliminary to demonstrating proof of concept and implementation in the marketplace:

- Energy markets exhibit non-linear, complex behaviors and structures.
- Traditional linear and partially linear forms of analysis are unable to identify complex structures on a consistent basis.
- Non-linear techniques (developed for the study of complex adaptive systems at the Los Alamos National Laboratory, the Santa Fe Institute, the University of Michigan, and elsewhere) which include the use of synthesis, genetic algorithms, neural nets, Monte Carlo simulations, Boolean logic, hill climbers, etc., can identify these complex structures on a more consistent basis.

Based on these assumptions, we took the following initial actions:

- Our first step was to identify the traditional linear and partially linear forms of analysis ("technical" marketing systems) typically used by energy traders.
- The most popular "technical" tools take the form of time series analyses that have popularized names such as Acceleration, Momentum, Volatility, Moving Average, etc.
- Though none of these tools in isolation can consistently identify complex structures, I hypothesized that a combination of these techniques through a voting mechanism (a hidden layer) might identify otherwise invisible complex structures. In other words, the interaction of these rules, like the interaction of traders on a trading floor, might mimic the complex behavior of the energy markets.

After considering the best approach, we took the following next steps:

- We analyzed the traditional marketing techniques and calculated a series of unique attributes that we used in defining optimum buy and sell rules for each technique. Our "new" rules were not the traditional rules of thumb usually associated with these technical marketing techniques.
- Among the unique attributes we calculated were:
    - Profit & loss (P&L) values at various buy and sell points for each technique;

    - Theoretical optimum for each day's marketing activities based on specific terminal values;
    - Total P&L for each rule compared with P&L per trade for each rule; and
    - Comparison of actual P&L versus theoretical optimum.
- Once these calculations were complete, we created graphs for each rule showing the increase or decrease of P&L across a range of entry points. We looked for P&L peaks that had gradual slopes and avoided sharp peaks (even if they indicated higher absolute P&L values). One of our goals was to minimize the value at risk, a technical measure of marketing risk, associated with the program.
- Though we investigated the use of a more complicated weighted-voting system (based on prior period results), we settled on a simple one virtual trader/one vote mechanism at the end of this initial study. Each technical marketing rule was polled each day, and the various buy and sell signals were netted to create a synthetic Black Box result. Separate rules were created for natural gas and crude oil.

Virtual Test Results (these results were calculated, but were not actually traded in the marketplace):

- Using the techniques described above and giving each vote the weight of a fixed number of standard futures contracts for natural gas and crude oil, the Black Box programs for both natural gas and crude oil were consistently profitable from July 1999 through June 2000.
- During the period July through December 1999, the natural gas Black Box accumulated hypothetical P&L of $6 MM, and the crude oil Black Box accumulated hypothetical P&L of $2 MM.
- During the period January through June 2000, the natural gas Black Box accumulated hypothetical P&L of $3 MM, and the crude oil Black Box accumulated hypothetical P&L of $2 MM.

During the summer of 1999, we looked at several alternatives suggested by scientists in Santa Fe for design of a second-generation version of the Black Box, to be called the "Pepper Box." We aimed to build a robust machine-aided learning system and predictive-pricing tool that could evolve and adapt to changing market conditions. The scientists in Santa Fe noted that many people had previously attempted to build automated "black box" type trading systems with limited success. Typically, these programs worked until they stopped working due to inevitable future regime changes in the marketplace – we wanted to avoid that pitfall. The second iteration of the Black Box, the Pepper Box, evolved in the following manner:

- Starting in the summer of 1999, we worked with scientists from Santa Fe on a next-generation complex, adaptive, predictive-pricing program known as the Pepper Box.
- The Pepper Box used the Black Box results as a starting point and included hidden layers to its synthesis.
- These additional layers attempted to identify more profitable combinations of the Black Box technical rules through a weighted voting mechanism. This mechanism employed thousands of "virtual traders" who combined the technical rules in all of their possible combinations and weightings.
- The features ultimately incorporated in the Pepper Box automated the Black Box technical rule generator/optimizer using updated pricing data, and we created a portfolio analyzer for companies following numerous Black Box or Pepper Box predictive-pricing programs.

The following is a synopsis of the various types of complex adaptive systems that we discussed building in 1999 – we ultimately settled on the "Pepper Box" approach:

- The Evolving Black Box: This version would have built an automated version of the Black Box but could also have bred combinations that resulted in over-weighting of successful rules. The rules are the foundation, but mutations could have occurred in the form of additions from the set of all possible studies. Additional fundamental rules could have been added. Again, the studies were the agents, and the genetic code would have still begun with a limited pool.
- The Rapid Monte Carlo: Based on current data, this version would have run every variation (or a representative proportion for continuous variables) of the study parameters for each rule individually. This program could have identified the most fit rule tuning and then reported recommendations, and this could have been implemented easily so long as only the study parameters were changed. If the trading levels were modified too, this would have become cumbersome. However, this approach accomplished most of the goals of the Evolving Black Box without the overhead because of the smaller search-space. Preliminary attempts at this version indicated the results did not change that much in total, among the rule tunings.
- The Pepper Box: This was the solution that we finally chose, and it was a true genetic algorithm where each agent is a trader that looked at several studies and created decision rules based on Boolean logic. (i.e., I trade on Momentum only if Acceleration

> confirms and Volatility also confirms). Virtual traders were then eliminated if their rules were consistently wrong. The most successful traders bred others, and mutations introduced new combinations of rules. The Pepper Box choice had a huge search space allowing for numerous rules and combinations. This also mimicked reality since no real trader would act solely on a single rule. In contrast to the other choices, this used the full power of the genetic algorithm approach.

IMPLEMENTATION OF THE BLACK BOX/PEPPER BOX PREDICTIVE PRICING PROGRAM

Based on the success of the Black Box program during the virtual testing period (July 1999 through June 2000), the Black Box/Pepper Box program was implemented on a limited trial basis to demonstrate proof of concept with actual transactions entered in the marketplace. Though the initial Black Box program ran on standard office computers, the additional complexity of the Pepper Box program required the use of a "Beowulf Cluster" of computers that provided a relatively low cost, high performance, parallel computing capability. Computing groups in Santa Fe ran the Pepper Box program each night for implementation the following day.

What we found almost immediately was the human traders' strong instinctive dislike and distrust for the often counter-intuitive recommendations provided by the Pepper Box program, and one of the first problems involved the periods of inactivity. Because both the Black Box and the Pepper Box were designed to identify long-term pricing trends and ignore short-term fluctuations, there were regular periods when no action was recommended, and this inactivity was particularly troubling to the human traders.

After struggling to explain that this inactivity was a feature and not a bug of the Pepper Box, we ultimately decided to introduce a small stochastic element (a form of Brownian motion or noise) into the program that yielded a few daily transactions. Based on our calculations, these small "daily" transactions would yield almost no net positive or negative results. But the reaction from the human traders was enthusiastic. They saw this revision as a huge improvement in the Pepper Box program since it gave them a daily activity to perform in the marketplace.

A much more significant problem in implementing the Pepper Box program involved the ongoing disagreements that the human traders had with the substantial long-term (as opposed to the stochastic "daily") Pepper Box recommendations. It was almost inevitable that some of the best and most useful Pepper Box predictions and proposed transactions were made at times when human traders either wanted to take no action or even wanted to take contrary actions. Explaining the logic of the Pepper Box transactions to the human traders was a hopeless task. The hidden layers, voting mechanisms, and weighting calculations made a simple explanation impossible.

The human traders were also reasonably concerned that the "fundamentals" of the marketplace (i.e., supply reports, demand changes, transport disruptions) were not explicitly entered into the Black Box/Pepper Box program. In fact, at the beginning of the project we debated whether to incorporate these traditional datasets. However, we decided after the initial study that the profound liquidity of the global crude oil and natural gas markets resulted in the implicit accounting for all these "fundamentals" in the rich pricing data that we fed into the Beowulf Cluster continuously. And the pricing data arrived in real time while the traditional fundamental datasets contained information that was often weeks or months in arrears.

Ultimately, though both the Black Box virtual marketing results and the Pepper Box actual marketing results were regularly profitable, we decided after a relatively short-lived experimental period to shut down this early effort to adopt and use primitive artificial intelligence technology. The primary role of the human traders within Occidental was to sell the crude oil and natural gas production of the company, and their marketing efforts were only a secondary activity (non-sales related marketing activity was authorized in part to keep the traders active in the marketplace even when not selling or hedging the company's oil and gas production). Since the company needed human traders for its essential product sales activity, and the human traders strongly disliked the Black Box/Pepper Box, we reluctantly shut it down.

CONCLUSION

The key to using artificial intelligence wisely and safely in the future will involve the successful integration of this new technology with its human partners. Like the complexity techniques that we used in designing the Black Box and Pepper Box programs (genetic algorithms and neural nets), explaining the internal logic of the new artificial intelligence programs will be almost impossible. Therefore, demonstrating the real and ongoing benefits understandably and convincingly will be critical to successful future adoption and use of artificial intelligence.

* * *