Expert Analysis

Artificial Intent: AI On The Trading Floor

By **Collin Starkweather and Izzy Nelken**
January 23, 2019, 1:21 PM EST

Collin Starkweather

Kara Stein, a departing member of the U.S. Securities and Exchange Commission, recently posed the question, “How do you regulate algorithms?”[1] The authors of a recent Harvard Business Review article advocated for the auditing of algorithmic decision-making and artificial intelligence, or AI, technologies, fearing potential “serious problems” associated with further introduction and development of those technologies.[2]

In the same vein, Microsoft recently began developing tools to help engineers identify unintended bias in algorithms, after machine learning, or ML, algorithms to predict whether defendants are likely to commit crimes in the future and facial recognition software were found to have “algorithmic bias.”[3]

Izzy Nelken

The fact that algorithms are being written to police other algorithms belies an important aspect of AI and ML that will likely have broader implications in a variety of areas, including regulation and compliance in financial markets. Recent developments in AI, ML and related technologies have seen developers of these technologies become ever farther removed from the results generated by these models. As the technologies continue to advance, the potential to draw inferences may be limited by the inherent nature of the technology.

This article will discuss some of these implications in the context of “spoofing,” one of a variety of trading activities that regulators consider to constitute illegal market manipulation. The Dodd-Frank Act defines spoofing as “bidding or offering with intent to cancel before execution.”[4] It has been the focus of increasing regulatory scrutiny in recent years, with several prominent cases involving spoofing either having recently been resolved or continuing to wind their way through the courts.[5] The question of intent not only has proven to be one of profound disagreement with respect to allegations of spoofing,[6] but is at issue with other kinds of trading practices of concern to regulators[7] as well as other matters of law more generally.

When the question arises in the context of trading practices, in tandem with the testimony of fact witnesses, the court may have the benefit of software design documents or a computer program (or code) from which to draw inferences.[8] However, with advances in and broader dissemination of ML, AI and related technologies, to the extent they find application in algorithmic trading, the potential to draw inferences even with access to code may be limited.

Spoofing serves as a good backdrop against which to consider the question of intent, not only because of the ready connection between trading strategies and implementations in software, and the court's focus on these algorithms as expressed in code, but also because of the rapidity with which leading-edge advances in technologies such as AI and ML are absorbed and deployed by trading operations.

The remainder of this article discusses those recent advances in technology which may inform the question of intent, not only with regards to spoofing, but more broadly in circumstances where the court may be disposed to draw inferences from code, as these technologies continue to advance and become more widely adopted.

The Evolution of Trading

While many traders remain "point and click" operators who manually enter trades, in much the same way as users of online retail brokerage platforms,[9] an increasing volume of trading is being performed algorithmically, independent of direct human intervention.[10]

In the first federal prosecution of spoofing in the U.S., Michael Coscia was indicted in October 2014 on six counts of commodities fraud and six counts of spoofing for executing a high-frequency trading strategy that involved entering orders "that he intended to immediately cancel before they could be filled by other traders."[11] During the trial, the prosecutor described Coscia's trading software as "designed to cancel" and "programmed to cancel."[12] At the time of Coscia's conviction, David Meister, the U.S. Commodity Futures Trading Commission's enforcement director, stated that "[w]hile forms of algorithmic trading are of course lawful, using a computer program that is written to spoof the market is illegal and will not be tolerated."[13]

To inform the question of intent, witnesses who have knowledge of the relevant facts and circumstances may provide testimony to assist the trier of fact. In their bid to establish intent, for example, prosecutors in the Coscia matter relied upon the testimony of Coscia's programmer and Coscia's own testimony regarding the design of the program he used for his high-frequency trading strategy.[14] Software design artifacts such as flow charts or use case diagrams may also serve to complement testimony regarding the communication of intent,[15] and the code that implements a trading strategy itself may even be analyzed.

When most people think of an algorithm, they think of simple, rule-based operations such as "if A then B." This mental model is consistent with the way trading algorithms have traditionally been designed.[16] Although it may be very complex, if the algorithm expresses a set of rule-based conditionals, given a set of market conditions, the outcome of the execution of the code will be unambiguous.

However, the technology available to algorithmic traders, and even the nature of what practitioners might consider to constitute an algorithm, has been evolving and becoming increasingly sophisticated with the development of tools and techniques to parse and analyze "big data" and the associated fields of machine learning and artificial intelligence.[17]

The Black Box

AI and ML do not describe singular, cohesive approaches to data processing and analysis. Rather, they serve as umbrella terms under which a wide variety of analytical techniques are categorized. While some of these techniques are rule-based, with any given set of inputs leading inexorably to a predictable output, other techniques have more of the character of a “black box.” They may produce desired results (such as the identification of a cat in a picture) but the precise mechanics of that determination may not be clear, even to those who implement the technique.[18]

The neural network is one such technique, for which even its own designers cannot explain precisely how given outcomes are produced.[19] Neural networks have been around since long before machine learning became a buzzword[20] and have been used for trading systems and financial analysis as far back as the 1990s,[21] but have recently seen a resurgence with the development of “deep learning” models.[22]

These systems rely on virtual neurons that lie between the input (such as a photo) and the output that is produced from the network (such as a determination of whether the photo contains a picture of a cat). The nature of the connections between those neurons is determined by a training dataset comprised of inputs for which the correct outputs are known.

Today, neural networks and deep learning technologies power facial recognition software used by Facebook and law enforcement authorities,[23] natural language processing used by digital assistants such as Amazon’s Alexa, Apple’s Siri and Google Assistant,[24] and in the identification of medical conditions such as skin cancer.[25] These technology firms benefit from massive repositories from which to draw training datasets for their algorithms. Facebook, for example, has billions of photos available in which users have tagged individuals that could potentially be used to train their facial recognition software.

Particularly with very complex models, it may not be easily discernible why a particular output is observed based on a particular input. For example, if a neural network identifies a cat in a photograph, it may not even be clear to the developer or developers who built and implemented the model exactly why the neural network produced a particular result. A reporter for the New York Times recently commented that “[machine learning] algorithms can’t articulate what they’re thinking. We don’t know why they work, so we don’t know if they can be trusted.”[26]

Within the domain of trading, ML has also found ready application in the field of surveillance and regulatory compliance. Surveillance tools had historically relied on detection of trading patterns using a set of rule-based criteria. In the case of spoofing, this might take the form of a set of conditionals based on the relative directionality, size and timing of orders.[27] Recently, trading surveillance technologies have adopted machine learning tools and techniques to identify patterns in trading activity such as spoofing much in the way that facial recognition software identifies an individual in a photograph.[28]

With a rule-based set of criteria, the reason that a certain pattern of trading activity has been identified is unambiguous. There is a “bright line” that distinguishes patterns, and those that fall on one side of the line are identified, while those that fall on the other are not, even if they otherwise may seem very similar.

However, if a pattern of activity has been identified by an AI- or ML-based model, as with image recognition, it may not be clear why that identification was made. Moreover, even identical implementations of a given model might produce different outcomes depending on factors such as the training dataset that was used.[29]

Looking Ahead

As AI and ML techniques are increasingly adopted for the purposes of trading, they introduce a layer of abstraction between the trader and the resulting trading activity, and the level of abstraction will only grow as the sophistication of these algorithms increases.

Looking even farther into the future, as automated trading becomes more sophisticated, regulators may even be faced with the prospect of trading software "discovering" unintended behaviors. Where would responsibility lie if, for example, a team of traders and data scientists implement an AI-based trading model that, independently of the intended design, engages in activity regulators identify as consistent with spoofing? Or if two or more such automated trading systems independently discover that they can profit from cooperating in a pattern of trading activity that would be identified by trading surveillance as spoofing if only one of the trading systems were to engage in the activity alone, but which is effectively hidden from surveillance when the systems cooperate?[30]

It's only a matter of time before trading systems become sufficiently sophisticated that these kinds of questions will need to be considered.

Collin Starkweather, Ph.D., is the founder of Starkweather Economics LLC, and is an economics and technology consultant working on complex litigation and other matters related to finance, competition economics and information technology.

Izzy Nelken, Ph.D., is the founder of Super Computer Consulting. He is a member of the CBOE Product Development Committee and has served as a consultant or expert witness on several matters involving allegations of market manipulation and other securities matters.

[1] Dave Michaels, "Machine Trading Needs More Oversight, Departing SEC Official Says," The Wall Street Journal, Dec. 21, 2018. In the article, she noted that securities market structures have "evolved into lots of computers talking to each other."

[2] James Guszcza, Iyad Rahwan, Will Bible, Manuel Cebrian and Vic Katyal, "Why We

Need to Audit Algorithms," Harvard Business Review, Nov. 28, 2018, available at https://hbr.org/2018/11/why-we-need-to-audit-algorithms. The article noted that the EU's General Data Protection Regulation requires that organizations be able to explain their algorithmic decisions, and that the city of New York recently created a task force to study possible biases in algorithmic decision systems.

[3] Kyle Wiggers, "Microsoft is developing a tool to help engineers catch bias in algorithms," Venture Beat, May 25, 2018, available at https://venturebeat.com/2018/05/25/microsoft-is-developing-a-tool-to-help-engineers-catch-bias-in-algorithms/.

[4] See Section 767.

[5] The U.S. Supreme Court recently declined to hear a certiorari petition by Michael Coscia, the first trader to be criminally prosecuted by the U.S. Department of Justice for spoofing. See, e.g., Greg Trotter, "Trader Michael Coscia 1st in nation to be sentenced under 'anti-spoofing' law," Chicago Tribune, July 13, 2016, available at http://www.chicagotribune.com/business/ct-spoofing-trial-sentencing-0714-biz-20160713-story.html; Jan Paul Miller, Steve Sherman and Amina Musa, "The Anti-Spoofing Provision of the Dodd-Frank Act: New White Collar Crime or 'Spoof' of a Law?" St. Louis Bar Journal, February 2016, available at https://www.thompsoncoburn.com/docs/default-source/News-Documents/spoofing.pdf; Ed Beeson, "Justices Won't Hear Challenge To Dodd-Frank 'Spoofing' Ban," Law360, May 14, 2018, available at https://www.law360.com/articles/1043291. In another prominent spoofing matter, UBS trader Andre Flotron was recently acquitted of spoofing. See, e.g., Christie Smythe, "Ex-UBS Metals Trader Beats Spoofing Conspiracy Charge," Bloomberg, April 25, 2018, available at https://www.bloomberg.com/news/articles/2018-04-25/ex-ubs-metals-trader-flotron-beats-spoofing-conspiracy-charge. Spoofing enforcement actions have been initiated in the last five years by the CFTC, SEC, FINRA and major exchanges. See, e.g., Trillium Surveyor, "Recent Trade Surveillance Enforcement Actions," available at https://www.trlm.com/knowledgebase/recent-trade-surveillance-enforcement-actions/.

[6] For example, in the opening statement of the Coscia trial, defense attorney Steven Peikin stated that one of the "two most critical issues in this case" was "did he place orders in the market with the intent to cancel them before they were traded," and on that topic "there is very, very serious disagreement." (Coscia Trial Transcript, October 26, 2015, at 173:1-6.)

[7] For example, CFTC Rule 575 defines prohibited "disruptive" trading practices as follows:

All orders must be entered for the purpose of executing bona fide transactions. Additionally, all non-actionable messages must be entered in good faith for legitimate purposes.

No person shall enter or cause to be entered an order with the intent, at the time of order entry, to cancel the order before execution or to modify the order to avoid execution;

No Person shall enter or cause to be entered an actionable or non-actionable message or messages with intent to mislead other market participants;

No Person shall enter or cause to be entered an actionable or non-actionable message or messages with intent to overload, delay, or disrupt the systems of the Exchange or other market participants; and

No person shall enter or cause to be entered an actionable or non-actionable message with intent to disrupt, or with reckless disregard for the adverse impact on, the orderly conduct of trading or the fair execution of transactions.

To the extent applicable, the provisions of this Rule apply to open outcry trading as well as electronic trading activity. Further, the provisions of this Rule apply to all market states, including the pre-opening period, the closing period and all trading sessions.

See,
e.g., https://www.cftc.gov/sites/default/files/filings/orgrules/14/08/rule082814cmedcm001.pdf
.

[8] The use of code in legal inference is evocative of Lawrence Lessig's maxim that "code is law," which Professor Lessig popularized in a book titled "Code and Other Laws of Cyberspace" published almost two decades ago. See, e.g., Lawrence Lessig, "Code is Law — On Liberty in Cyberspace," Harvard Magazine, Jan. 1, 2000, available at https://www.harvardmagazine.com/2000/01/code-is-law-html.

[9] Although the idea is the same, professional trading platforms (such as those offered by Interactive Brokers or Trading Technologies) offer a more sophisticated trading environment that that offered by online brokerages targeting retail investors (such as those offered by E*Trade or TD Ameritrade).

[10] For example, the U.S. Congressional Research Service estimated that high-frequency trading, a prominent subcategory of algorithmic trading, has increased substantially in the last decade or so to constitute around 55 percent of trading in U.S. equities and between two-thirds and four-fifths of trading in U.S. foreign exchange, interest rate and Treasury futures as of 2014-2015. (Rena S. Miller and Gary Shorter, "High Frequency Trading: Overview of Recent Developments," Congressional Research Service, April 4, 2016, available at https://fas.org/sgp/crs/misc/R44443.pdf.)

[11] "High-Frequency Trader Indicted for Manipulating Commodities Futures Market In First Federal Prosecution for 'Spoofing'," U.S Department of Justice, Northern District of Illinois, Oct. 2, 2014, available at https://www.justice.gov/usao-ndil/pr/high-frequency-trader-indicted-manipulating-commodities-futures-markets-first-federal.

[12] Coscia Trial Transcript, Oct. 29, 2015, at 847:2; Coscia Trial Transcript, Nov. 3, 2015, at 1421:10-11.

[13] CFTC Release Number 6649-13, “CFTC Orders Panther Energy Trading LLC and its Principal Michael J. Coscia to Pay $2.8 Million and Bans Them from Trading for One Year, for Spoofing in Numerous Commodity Futures Contracts,” July 22, 2013, available at https://www.cftc.gov/PressRoom/PressReleases/pr6649-13.

[14] Coscia Trial Transcript, Oct. 26, 2015, at 162:10-13 Coscia Trial Transcript, Oct. 29, 2015, at 846:25-847:7.

[15] In sophisticated software architecture and development exercise, for example, Unified Modeling Language artifacts such as use case diagrams, class diagrams, or activity diagrams may be generated.

[16] For example, Coscia described his trading instructions to his programmer in terms of rule-based “if/then” operations such as, “If the market is two tick[]s wide, bid for 10 contracts, then bid for 10 contracts. If it's less than 10 contracts, don't bid.” (Coscia Trial Transcript, Oct. 27, 2015, at 446:6-17.)

[17] Machine learning is sometimes used synonymously with the term “data mining,” which involves the “extraction of implicit, previously unknown, and potentially useful information from data,” though it properly refers to the set of tools and techniques used to accomplish the broader task of data mining. (Ian H. Witten and Eibe Frank, Data Mining: Practical Machine Learning Tools and Techniques, Second Edition, 2005, p. xxiii.)

[18] See, for example, Will Knight, “The Dark Secret at the Heart of AI,” Technology Review, April 11, 2017, available at https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/.

[19] Note that although neural networks are used as an exemplar in this article for the purposes of exposition, other such techniques, such as Bayesian clustering and networks, are widely used in machine learning and artificial intelligence applications.

[20] Important early work in establishing the theory underlying artificial neural networks was done as far back as the 1940s. See, for example, N. Yadav et al., An Introduction to Neural Network Methods for Differential Equations, Springer Briefs in Computational Intelligence, p. 13, available at https://www.springer.com/cda/content/document/cda_downloaddocument/9789401798150-c2.pdf.

[21] See, e.g., Jayesh Bapu Ahire, “Real world Applications of Artificial Neural Networks,” Medium, April 9, 2018, available at https://medium.com/@jayeshbahire/real-world-applications-of-artificial-neural-networks-a6a6bc17ad6a; Justin Sirignano, “Deep Learning Models in Finance,” SIAM News, June 1, 2017, available at https://sinews.siam.org/Details-Page/deep-learning-models-in-finance-2.

[22] In their seminal book on machine learning, Witten et al. observe that "[d]eep learning has sparked a renaissance of neural network research and applications." (Ian H. Witten, Eibe Frank, Mark A. Hall and Christopher J. Pal, Data Mining, Fourth Edition (2017), p. 419.)

[23] See, for example, Yaniv Taigman, Ming Yang, Marc'Aurelio Ranzato, Lior Wolf, "DeepFace: Closing the Gap to Human-Level Performance in Face Verification," Conference on Computer Vision and Pattern Recognition (CVPR), June 24, 2014, available at https://research.fb.com/publications/deepface-closing-the-gap-to-human-level-performance-in-face-verification/; Drew Harwell, "Facial recognition may be coming to a police body camera near you," Washington Post, April 26, 2018, available at https://www.washingtonpost.com/news/the-switch/wp/2018/04/26/facial-recognition-may-be-coming-to-a-police-body-camera-near-you/; Simon Denyer, "Beijing bets on facial recognition in a big drive for total surveillance," Washington Post, Jan. 7, 2018, available at https://www.washingtonpost.com/news/world/wp/2018/01/07/feature/in-china-facial-recognition-is-sharp-end-of-a-drive-for-total-surveillance/.

[24] See, for example, Steve Ranger, "Amazon Echo: The four hard problems Amazon had to solve to make it work," ZDNet, Sept. 14, 2016, available at https://www.zdnet.com/article/amazon-echo-the-four-hard-problems-amazon-had-to-solve-to-make-it-work/; Larry Dignan, "Google Assistant auditions to be your personal digital twin via Duplex," ZDNet, May 8, 2018, available at https://www.zdnet.com/article/google-auditions-to-be-your-personal-digital-twin-via-duplex-google-assistant/.

[25] H. A. Haenssle et al., "Man against machine: diagnostic performance of a deep learning convolutional neural network for dermoscopic melanoma recognition in comparison to 58 dermatologists," Annals of Oncology, May 28, 2018, available at https://academic.oup.com/annonc/advance-article/doi/10.1093/annonc/mdy166/5004443.

[26] Steven Storgatz, "One Giant Step for a Chess-Playing Machine", The New York Times, Dec. 26, 2018, available at https://www.nytimes.com/2018/12/26/science/chess-artificial-intelligence.html. The article concerned AlphaZero, an AI implementation by DeepMind, a division of Google, to play chess. In describing its style of play, Demis Hassabis, the founder and CEO of DeepMind, said that AlphaZero "doesn't play like a human, and it doesn't play like a program. It plays in a third, almost alien, way." Will Knight, "Alpha Zero's 'Alien' Chess Shows the Power, and the Peculiarity, of AI," Technology Review, Dec. 8, 2017, available at https://www.technologyreview.com/s/609736/alpha-zeros-alien-chess-shows-the-power-and-the-peculiarity-of-ai. Consistent with that observation, the author of the New York Times article further commented that "AlphaZero gives every appearance of having discovered some important principles about chess, but it can't share that understanding with us."

[27] See, for example, Jeffrey A. Brown, Steven Pellechi, Thomas Cordova and Tanner Kroeger, "Department of Justice Broadens Aim on Spoofing Enforcement," Lexology, May 4, 2018, available at https://www.lexology.com/library/detail.aspx?g=419b7c97-0d54-4633-

90bf-169b7ca2bcec.

[28] For example, Trading Technologies and NASDAQ have incorporated machine learning into their market surveillance platforms. (Jay Biondo and Morgan Trinkaus, Product Managers, Surveillance, “Trade Talk Blog: Make Surveillance Smarter and Eliminate Blind Spots With TT Score,” Jan. 18, 2018, available at https://www.tradingtechnologies.com/blog/2018/01/18/make-surveillance-smarter-and-eliminate-blind-spots-with-tt-score/; Bob Violino, “Nasdaq invests in artificial intelligence,” ZDNet, Mary 24, 2017, available at https://www.zdnet.com/article/nasdaq-invests-in-artificial-intelligence/.

[29] If, for example, the training dataset were comprised of trading activity identified by regulators as spoofing, the training dataset would expand over time as additional trading activity is identified by regulators.

[30] While the idea that AI implementations may “discover” unorthodox and unexpected ways to cooperatively achieve objectives may seem far-fetched, it has recently been reported that AI implementations on Facebook assigned the task of bartering actually created their own language for the purpose of the bartering exercise. Andrew Griffin, “Facebook’s Artificial Intelligence Robots Shut Down After They Start Talking to Each Other in Their Own Language,” The Independent, July 31, 2017, available at https://www.independent.co.uk/life-style/gadgets-and-tech/news/facebook-artificial-intelligence-ai-chatbot-new-language-research-openai-google-a7869706.html.