accenture

Strategy | Consulting | Digital | Technology | Operations

Accenture Labs

Ethical algorithms for "sense and respond" systems

High performance. Delivered.

Introduction

Sense and respond (S&R) systems have been in use for decades. Responding to their environment in real time with little or no human input, these systems have been used successfully across a diverse range of applications, including lunar landers, anti-lock brakes, and automatic financial trading systems.

As S&R systems become more pervasive, it is crucial that they be designed thoughtfully and ethically so that they improve how we work and live, not increase risk or the potential for harm to individuals or society at large.

These developments are being driven by recent advances in computing and technology, which have enabled S&R systems to handle complex decision-making in real-time. At the same time, machine learning allows algorithms to learn from data and react to new scenarios never explicitly encountered and/or considered by their creators. As the prevalence and decision-making capabilities of S&R systems continue to increase, there's great potential for ethical failures with significant impacts. Organizations developing S&R systems must carefully consider the ethical contexts of their systems and how they should address these new forms of risks. If ethics are not properly considered during S&R design, implementation, and use, they can propagate unwanted biases and erode human trust in both the systems and the organizations that deploy them. Practitioners designing S&R systems that will become a key part of consumer life must take these risks into account.

This paper focuses on physical S&R systems operating in "everyday life" environments (rather than closed or controlled environments like factories) and is primarily concerned with systems that respond to stimuli in real time. It offers foundational tenets for designing ethical S&R algorithms, and recommendations for actions that practitioners can take to address the ethical concerns that these systems raise.

Ethical issues

Understanding the ethical issues around S&R systems begins with understanding and acknowledging biases found in machine-learning algorithms. The difficulty in handling these biases and grappling with machine decision-making can be demonstrated with the trolley problem thought experiment and by comparing human evaluations of human-made and machine-made decisions.[1] These differences reveal themselves when S&R systems are considered for applications where human autonomy is removed, as well as when the objective for decision-making emphasizes an entire network rather than an individual actor in a system.

Ethical issues in machine-learning systems

Sense and respond systems are a type of machine-learning (ML) system, and are subject to the array of ethical concerns that apply to all ML systems. These issues can arise due to biased and flawed human decisions in the process of designing the system or in the interpretation of the system's output. There are also classes of ethical issues that exist in the data and in the statistical models of that data.

Figure 1: Entry points for bias in data-science processes for sense and respond systems

System design	Modeling and training	Presentation and implementation
• Human cognitive biases • Algorithmic aversion • Algorithm selection • Data collection bias • Missing or misquantified inputs	• Reinforcement bias • Societal bias • Safety boundaries • Fairness and under-representation of minority classes • Validation of data labels	• Flawed interpretation of results • Misapplication of models • Disregarded design assumptions • Verification of data labels

Data collection bias, for example, can be demonstrated with the well-known example of Boston-based StreetBump, which developed a smartphone app that used GPS and accelerometer data to identify and report potholes in the city's streets. This was a well-meaning attempt to allocate city resources to patch thousands of potholes. Collecting data via a smartphone app, however, failed to take into account that many of the city's elderly and lower-income residents do not have smartphones. The designed data collection approach generated data sets biased in favor of wealthy neighborhoods, and neglected those with fewer resources. The city eventually took steps to control for this bias by statistical means, but it remains an illustrative example of unintended bias in data collection.[2]

In the design process for a machine-learning system, where engineers architect the system and its data flows, human cognitive biases and algorithmic aversion are potential concerns.[3] These issues may be compounded during the training phase, in which the algorithm learns from the available data. Biased or incomplete data can build flawed statistical models and reinforce societal biases.

Flawed interpretation of the output of machine-learning systems leads to another class of failures. This can occur via a number of missteps: the inability to acknowledge the limitations of a machine-learning system; presentation of output without context or disclosure of design assumptions; or the use of trained models in applications other than those intended (see Figure 1).

These ethical issues are relevant and critical to all machine-learning systems and data products, and should be considered carefully by creators and users of such systems. These issues are often even more critical in sense and respond systems, which interact directly with humans and the physical world.

The "Trolley Problem"

There are specific ethical issues relevant or even unique to sense and respond systems. One thought experiment that evokes the quintessential ethical concerns faced by S&R systems is the famous "Trolley Problem." Traditionally, the problem is presented like this: a trolley is traveling on a track to which five people are tied; the trolley's path can be switched to a different track, on which a single person is tied. By doing nothing, five people will die; by making a decision to change course, only one person will die. Should an observer, standing next to the lever that would switch the trolley to the other track, change the trolley's path?

With decisions such as these now placed in the hands of machines rather than humans, as with automated public transit and self-driving cars, the questions raised in the Trolley Problem are increasingly relevant across a range of S&R systems. A reasonable assumption is that an algorithm will adopt a utilitarian perspective, weighing five lives against one life and concludi that switching paths is the morally correct choice because it maximizes the greatest good for the greatest number of people. But questions raised by the Trolley Problem challenge the utilitarian position: should a self-driving car ever sacrifice its passengers to prevent a larger accident? Can an algorithm really prevent a worse accident with a high enough confidence level? What is its moral responsibility or obligation when considering such a scenario? If forced to compare human lives, what should an algorithm take into consideration – are passengers more valuable than bystanders, should saving children be the highest priority? How location-dependent are these choices (e.g. would an autonomous vehicle make a different decision in Germany versus the United States)?

The utilitarian perspective is further limited by context and available data. If lives are commensurable, what data inputs should designers provide to the algorithm so that it can draw a comparison? What biases will authors inadvertently bake into their systems? The "greatest good" cannot be maximized if it cannot be objectively measured.

With the development of self-driving vehicles, the Trolley Problem moves from a hypothetical thought experiment to an ethical consideration of immediate relevance. More broadly, organizations dealing with these systems should understand the gravity of these design choices. While spending time designing for every hypothetical is counterproductive, acknowledging that these risks exist and understanding realistic approaches for addressing them fairly is key. This begins with considering the differences between human and machine decision-making.

Machine decisions vs. human decisions

Humans trust smartphones to track calendar events, or connected thermostats to control the temperature in their homes. Yet humans evaluate autonomous decisions differently than those made by other humans, and many would hesitate to accept a machine decision that had a direct and major implication for human life. Research shows that humans are more forgiving of mistakes made by other humans, and more critical of decisions made by algorithms.[4]

There are two fundamental reasons for this biased response. First, humans do not empathize with machines, which are non-living, non-thinking systems—and vice-versa. In a judge's decision to condemn an individual to life in prison, the context of that human's life and their state of mind at the time of the offending action is considered relevant. These considerations, however miniscule, are difficult at best for a machine to consider. Second, the machine "thought process" is generally considered cold and callous. By contrast, for humans, ethical decisions require empathy, an understanding of human values, and the ability to evaluate decisions in multiple contexts simultaneously. If humans do not believe machines are capable of such cognition, then they would naturally be mistrustful of a machine's decision in a case with ethical implications, or worse, a case that decided the future of another human being's life.

Autonomy

Cognitive psychology has increased our awareness of the many ways in which ideal rational cognition in humans can be compromised, and the great influence that non-conscious processes have in our decision-making. This awareness has created an active ethical discussion about when it's reasonable to override an individual's autonomy for their own benefit, and to help them arrive at meaningful decisions. This discussion is particularly familiar in a medical context: when should a caregiver override a patient's freedom of choice or prior directive?

Consider an injured human who might decide to forego medical treatment that would extend their life, but with a lower quality of life than before. If this decision is being made with knowledge of all appropriate facts, then it's the medical practitioner's duty to respect their right to choose. However, making sure that the decision-maker has a full understanding of the implications of their actions is not a straightforward task. In order for one to act autonomously, different decision-making criteria must be presented objectively without influencing bias. Conventional assumptions around autonomous decision-making are challenged by sense and respond systems that remove agency from an individual.[5]

The increasing prevalence of sense and respond systems suggests an analogous question: in what conditions would an autonomous system have the moral right to override a human choice or action?

Removing autonomy from an individual aligns with the fact that their practical autonomy may be non-existent; if stripping agency from individuals becomes necessary, it will not be taken lightly, particularly in societies where individual liberty is highly valued.[6]

As S&R systems become more ubiquitous, what effort must be put toward respecting individuals' ability to make their own decisions? What's more, when S&R systems do not use all available information and make poor choices, is the user or the designer at fault? Informed consent agreements that lay out what authority is being delegated to the S&R system, in terms that a user can easily understand, could help mitigate these risk vectors by eliminating the ambiguity around responsibility and authority.

"Network" vs. "Node" thinking

In cases where decision-making responsibility is handled by S&R systems rather than humans, there is added potential for larger-scale impacts. While an individual human is limited in the information that is both available to them and that can be considered for any one decision, a highly-connected algorithmic system may have vast sources of inputs as well as the computational power to consider those inputs simultaneously. As a result, it can perform decision-making that optimizes results on a systems level, rather than at the individual level.

Imagine an energy management system tasked with maintaining a comfortable temperature. During an extreme hot weather event, if the energy demands of maintaining an ideal indoor air temperature in a set of large buildings would likely cause a widespread blackout, the system might decide to accept a higher indoor air temperature in order to keep the grid operational. But such a decision would be possible only if the system had been designed with decision criteria that included reliable grid operation and if it possessed the necessary input data to recognize such a situation. Following such a decision, a key challenge is to convey how the system's decision reflects the best interests of the network as a whole and provides for more sustainable operations. Transparency into decisions can create a balance between the long-term consequences of decisions about networks and the short-term consequences that impact individual nodes.[7]

In designing inputs into a system, and the criteria for which to optimize, designers are making critical decisions about the way humans will engage with the world. These selections define the context of the system's decisions, and are a key area of consideration when addressing ethical issues. This is why appropriate and holistic planning is so critical. Even in this algorithm design phase, human nature can work against longer-term goals because humans struggle when considering an immediate payoff over delayed rewards, even when a slight delay yields a much greater reward.[8]

The network-node choice becomes increasingly important as systems scale. Building on the energy example from above, a 2012 report by the American Council for an Energy Efficient Economy found that the United States' overall energy consumption could be reduced between 12 and 22 percent by capitalizing on the data produced by existing infrastructure. S&R systems that help power companies manage outages, expected system loads, and the current state of their equipment could all be contributing factors to this potential reduction.

Recommendations

These ethical challenges in the design and use of sense and respond systems must be given due consideration in order to realize their full benefit to society. By acknowledging ethical concerns, designers and users can ensure successful long-term adoption, and avoid significant negative consequences for themselves and the communities they hope to serve. There are specific and tangible steps that can be taken to address these ethical constraints and mitigate their associated risks.

Auditability, transparency, and recourse

Designers must understand that there are many ways a sense and respond system could cause harm or commit ethics blunders. The critical route to success requires a strategy not just for prevention, but also for early recognition. This strategy mus also encompass public, transparent responses to failures and the ability for subjects of these failures to have recourse, with a direct feedback loop to the system designers and/or operators.

One component of such a strategy is to create roles and processes around auditability. Just as an ethics review board may be called upon to evaluate human subject research, similar specialized task forces should review ethics failures in automated systems. This review process should be open and transparent to the extent possible within the constraints of business and legal limitations.

The process for monitoring and analyzing the behavior of autonomous systems differs from the human analog in several important ways. First, it's possible to incorporate auditability directly into the system, capturing inputs and steps throughout the decision-making process. This can be automated through the use of sensors on the edge of systems that interface with the physical world, but it could also take the form of human-driven recourse when anomalies are detected. By embedding these capabilities into a system, a review board can recreate the system's internal representation of the scenario, including the inputs it was using and the relevant decision algorithms. This review process can result in algorithms being updated to better account for ethical decision-making.

Additionally, organizations should attempt to collect user feedback on the ethical concerns identified as relevant to their system. This feedback should be incorporated into an established internal auditing process, or maintained as a separate measure of success or failure. S&R systems interact with the physical world in a direct way, and all humans that witness or participate in that interaction can provide critical perspectives safely removed from the design and implementation process. Their continuous feedback is also important to properly monitor the changing behavior of the system over time.

At a minimum, users of S&R systems should be given clear paths for recourse when anomalies occur. It will be a long time before machines are able to compete with humans on anomaly detection. System designers should use this highly evolved skill to their advantage. But they will only be able to do so if they provide a simple and intuitive way to solicit that feedback.

It is likely that public reaction to the first major failures of autonomous vehicles (e.g. accidents involving serious human injury or fatalities) will include outrage and anger disproportionate to the failure rate of such systems – even though it will almost certainly be far below that of human drivers. Yet, organizations should publicly acknowledge mistakes when they occur, alleviating the risks from even the appearance of dishonesty or unethical behavior. Over time, transparency around failures, and the resulting changes in processes and algorithms, will be important to advancing consumer trust and acceptance of sense and response systems.

Google's monthly self-driving vehicle reports document and catalog road accidents in which their vehicles have been involved. In February 2016, the first instance of a Google self-driving vehicle causing an accident was reported. That month's report walks the reader through the accident, revealing what assumptions were incorrectly made by the car, as well as the steps taken to prevent future errors from happening.[9] Google's swift response and transparent handling of this situation is a good example of how to appropriately handle scenarios where sense and respond systems unpredictably cause harm.

Training and best practices

Organizational best practices reach beyond external transparency. They should also include processes to make system designers more aware of the unwanted biases that S&R systems can learn and present. This topic is covered in-depth in the companion piece in this series, Facilitating ethical decisions throughout the data supply chain.

S&R project leaders should ensure, through education and training, that all team members have at least a basic understanding of the ethical considerations relevant to the project. This requires that leadership, together with informed and independent input, identify such considerations at the beginning of any project. Periodic reviews of issues throughout the project, particularly at major milestones during algorithm training and testing, should all have ethics review boards or similar governance structures in place. In the same way that project managers are expected to identify and mitigate other types of project risks, ethics risk and mitigation should be considered critical.

Training for project managers and data scientists should strive to increase awareness of the ways that human discrimination and bias may be introduced into S&R systems, whether from individuals themselves or from the bias that is inherent to the use of data in certain contexts.

Specific techniques and tools should be identified that may assist in detecting discrimination, whether intended or otherwise, so that S&R system designers can address these issues.[10] Team managers should incorporate ethical considerations into existing "best practices" or design guidelines. Tramér, et al, suggest that algorithmic biases should be treated like bugs in code: such treatment encourages the mentality within a team that biases are detectable, that they can be fixed, and that, like any other bug, their presence affects product release.[11]

Similarly, as with software projects, it's possible to simulate high-risk scenarios to identify ethical concerns, and then provide specific instructions to the system to address them where necessary. Repeated simulation can be performed using subsets of data, as suggested by Suresh Venkatasubramanian.[12] Such simulations could determine whether the same results are produced for different populations or scenarios, or whether the system may be responding in a biased manner as a result of inadvertent bias on the part of its human designers.
A 2011 study from the University of Texas at Dallas, for example, found significant differences between east Asian and western facial recognition algorithms.[13] Algorithms were shown to perform better at identifying faces of individuals from their respective geographic region. With additional training focused on a broader set of human features, it's conceivable that the algorithms could have performed better. Such tools and techniques for a priori identification of ethical issues should be used wherever possible.

Regulation and policy

Demonstrable self-regulation highlights an organization's commitment to transparency and best practices. By participating in such efforts, organizations position themselves to improve the usefulness of new industry and legal policies and regulations being introduced to mitigate substantial related risks.

Consider an example from the financial industry. The risks from a single financial trading algorithm may seem insignificant, but a 2014 report from the SEC estimated high-frequency trading to represent at least 50 percent of the total volume of US-based trades.[14] At such scale, risks change fundamentally and new implications can emerge. With algorithmic trading, however, the scale does not even need to be that big: in 2010, a single trader, at his home in London, is alleged to have been responsible for the flash crash that caused nearly US$1 trillion to be erased from the US stock market.[15] Although this was not his intention, the outcome was severe. Ethical risks from machine-learning systems may also scale in unexpected ways: for instance, if everyone uses the same search algorithm, or if self-driving cars become commonplace. To what extent can the authors of these algorithms be expected to self-regulate the ways in which they may collectively weaken a larger system through their participation? Regulation and policy must play a role in governing these algorithmic ecosystems and mitigating such risks before S&R systems reach a critical mass.

Better regulation and policy around S&R systems can help by outlining better safeguards against harm. In doing so, regulators can force organizations to consider the implications of their algorithmic actions and hold them accountable for ethics failures in the implementation of these algorithms.

In this context, consider the US stock market's flash crash, caused by algorithmic trading. In an effort to prevent future systemic failure, the SEC put "circuit breakers" in place that would detect if certain stocks were being traded under conditions similar to the 2010 crash, and halt trading for five minutes to allow imbalances to work themselves out. On August 24, 2015, these breakers were tripped 1,200 times throughout the trading day. Traders were acting legally, but without the breakers, their collective actions could have resulted in a significant failure that would have reduced public trust in the stock market.

Moving past the risk of individual ethics failures, the "circuit breaker" example illustrates another risk: while these systems may have prevented a larger crash, their "trips" stacked on one another and had the unintended result of preventing stock prices from rebounding in a reasonable amount of time.[16,17] Even if fail-safes fulfil their original purpose, certain cases may yield results where the fail-safes work too well and over-react. Overly cautious behavior around an S&R system may initially reduce benefits; however, these systems must be iterative in order to succeed. Organizations will need to learn to adjust sensitivities and fine-tune them for specific systems.

Conclusions

Sense and response systems offer tremendous potential and are becoming increasingly prevalent in human environments. Yet, being designed and implemented by humans, these systems are subject to the same types of bias and flawed decision-making that affects individual human cognition. As a result, they pose considerable ethical concerns, which their human designers have a moral imperative to address.

Given the potential to set precedent, practitioners today can begin to increase their efforts around auditability and transparency in order to improve visibility into the design process and how ethics failures are addressed. This kind of behavior is encouraged if system designers are trained to recognize ethical issues and to have best practices in place for addressing unwanted bias from input variables. When organizations take strides to produce S&R systems with an understanding of the ethics risks that a system could pose, new regulatory regimes will not result in a decreased capacity to innovate.

Much can be learned from considering the ethics issues associated with other machine-learning systems. But S&R systems offer their own unique challenges and, as they increasingly operate in real-time physical environments with humans, the consequences of ethics mistakes may be greater than in other machine-learning contexts.

Today's designers of S&R systems will lay the foundations upon which future designers will build. This precedent will set the standard for best practices, and will determine whether or not S&R systems will earn the public's trust, which is critical to ensuring widespread adoption and use of these systems.

Establishing trust will not be a simple or quickly achieved endeavor. Major challenges must be overcome: the fear of humans being replaced, and even controlled, by emotionless, autonomous systems whose operation is largely a mystery to the public, and the idea that a machine, no matter how well designed, could never "think" about a decision the way that a human could—at least for many decades to come.

System designers must address these and other challenges by embracing the long game of transparency, acknowledging lapses, and providing mechanisms for recourse. Additionally, designers must demonstrate a willingness to respond when these systems make mistakes or when humans make mistakes in their design. These steps must be considered from day one of any S&R system design and continuously reinforced throughout its development and use.

With careful planning and a demonstrable interest in addressing the ethical concerns associated with these technologies, S&R systems can be put in place with the best possible opportunity to gain human trust. These systems offer tremendous potential benefits to society—including great improvements to human safety —that can only be realized by establishing the public's trust. Trust must come first.

Minimizing risk exposure from data: 100-day/365-day recommendations

In the next 100 days, try to understand how your business interacts with external sense and respond systems and how other businesses interact with the sense and respond systems you provide. Understanding the relationships within and between S&R systems will help your organization to develop a plan for their best practice use:

- Document the processes your organization takes to address ethics risks and failures in sense and respond systems. For each step of each process, identify how that action could improve or degrade customer trust.
- Identify and make a list of currently deployed sense and respond systems. Highlight the ones where human or algorithmic bias may have played a role in determining outcomes. Do the same for systems under development.
- Catalog any external sense and respond systems your organization engages with and enumerate the potential biases from Figure 1 that could affect the organization or its customers.
- Create a list of potential circuit-breakers to prevent your sense and respond systems from being too far outside of specification, placing the ones that mitigate the most risk at the top of the list.
- Challenge your data scientists to use FairTest features integrated in SciPy to discover any unintended and unwarranted biases in your algorithms.

In the next 365 days, your organization should have a better understanding of the S&R systems with which it interfaces. Evaluating these connections through an ethical lens will help with decisions around increasing or decreasing the number of S&R systems you currently use:

Leverage the catalog of potential biases introduced from external sense and respond systems to test for each potential bias. Collaborate with vendors wherever possible and share results with them.

Using the list of circuit-breakers, implement at least half of the top-ten circuit-breakers identified. Keep a dashboard of how frequently the circuit-breakers are invoked and use that to improve your algorithms and systems.

Work with your marketing, alliance, public relations, and product management teams to maximize transparency with customers and stakeholders about any flaws discovered during your various ethics investigations and tests. Be sure to include the conditions which cause biased outcomes, what you're doing to improve the system, and what customers can expect moving forward.

Deploy an internal award and recognition program for employees who identify ethics risks with internal and customer-facing systems. Invite those who surface issues to join the engineering and data science conversations on how to address the issue.

Contact Us

Steven C. Tiell
Senior Principal, Digital Ethics
Accenture Labs
steven.c.tiell@accenture.com

Aman Ahuja
The Data Guild
aman@thedataguild.com

Harrison Lynch
Senior Research Analyst,
Digital Ethics
Accenture Labs
harrison.lynch@accenture.com

Michelle Sipics
Principal, Technology Vision
michelle.sipics@accenture.com

Contributors

Cameron Turner
The Data Guild
cameron@thedataguild.com

Scott L. David
University of Washington
sldavid@uw.edu

Gideon Wulfsohn
Tufts University
gideon.wulfsohn@tufts.edu

References

[1] Trolley problem. (2016, April 7). In Wikipedia, the free encyclopedia. Retrieved from https://en.wikipedia.org/w/index.php?title=Trolley_problem&oldid=728821194

[2] Crawford K (2013) The Hidden Biases in Big Data. Harvard Business Review. Available from: https://hbr.org/2013/04/the-hidden-biases-in-big-data (accessed 24 April 2016).

[3] Dievorst B, Massey C, and Simmons J (2014) Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err. Journal of Experimental Psychology: General 144(1): 114–126.

[4] Dievorst B, Massey C, and Simmons J (2014) Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err. Journal of Experimental Psychology: General 144(1): 114–126.

[5] Coekelbergh M (2015) Care Robots and the Future of ICT-mediated Elderly Care: A Response to Doom Scenarios. AI & Society: 1-8.

[6] Butkus, M. A. (2015). Free Will and Autonomous Medical Decision-Making. Details Journal of Cognition and Neuroethics (ISSN: 2166-5087). March, 2015, 3(1), 75–119.

[7] Rose A, Vadari S, and Wigle L (2014) How the Internet of Things Will Enable Vast New Levels of Efficiency. ACEEE Summer Study on Energy Efficiency in Buildings 9: 295-307.

[8] Remington B, Warry C, and Sonuga-Barke E (1999) When More Means Less: Factors Affecting Human Self-Control in a Local versus Global Choice Paradigm. Learning and Motivation Vol 30(1): 53-73.

[9] Google Self-Driving Car Project Monthly Report February 2016. Google. Available from: https://static.googleusercontent.com/media/www.google.com/en//selfdrivingcar/files/reports/report-0216.pdf (accessed 24 April 2016).

[10] These types of tools are starting to emerge; one example is FairTest from a team at Columbia. This tool can be used to test algorithms for bias on a set of "protected variables" (race, age, gender...). See http://arxiv.org/abs/1510.02377

[11] Tramer, F., Atlidakis, V., Geambasu, R., Hsu, D., Hubaux, J.-P., Humbert, M., ... Lin, H. (2015). Discovering Unwarranted Associations in Data-Driven Applications with the FairTest Testing Toolkit. arXiv Preprint arXiv:1510.02377. Retrieved from http://arxiv.org/abs/1510.02377

[12] Intagliata C (2015) Why Machines Discriminate - and How to Fix Them. NPR. Available from: http://www.sciencefriday.com/segments/why-machines-discriminate-and-how-to-fix-them/ (accessed 24 April 2016).

[13] Phillips P et al. (2011) An other-race effect for face recognition algorithms. ACM Transactions on Applied Perception Vol 8(2): article 14.

[14] (2014) Equity Market Structure Literature Review Part II: High Frequency Trading. U.S. Securities and Exchange Commission. Available from: https://www.sec.gov/marketstructure/research/hft_lit_review_march_2014.pdf (accessed 24 April 2016).

[15] Levine, M. (2015, April 21). Guy Trading at Home Caused the Flash Crash. Retrieved July 8, 2016, from https://www.bloomberg.com/view/articles/2015-04-21/guy-trading-at-home-caused-the-flash-crash

[16] Egan M. (2015) Trading was halted 1,200 times Monday. CNN. Available from: http://money.cnn.com/2015/08/24/investing/stocks-markets-selloff-circuit-breakers-1200-times/index.html (accessed 24 April 2016).

[17] (2010) Findings Regarding the Market Events of May 6, 2010. U.S. Securities and Exchange Commission. Available from: https://www.sec.gov/news/studies/2010/marketevents-report.pdf (accessed 24 April 2016).

About Accenture Labs

Accenture Labs invents the future for Accenture, our clients and the market. Focused on solving critical business problems with advanced technology, Accenture Labs brings fresh insights and innovations to our clients, helping them capitalize on dramatic changes in technology, business and society. Our dedicated team of technologists and researchers work with leaders across the company to invest in, incubate and deliver breakthrough ideas and solutions that help our clients create new sources of business advantage.

Accenture Labs is located in seven key research hubs around the world: Silicon Valley, CA; Sophia Antipolis, France; Arlington, Virginia; Beijing, China; Bangalore, India, Dublin, Ireland, and Tel Aviv, Israel. The Labs collaborates extensively with Accenture's network of nearly 400 innovation centers, studios and centers of excellence located in 92 cities and 35 countries globally to deliver cutting-edge research, insights and solutions to clients where they operate and live. For more information, please visit www.accenture.com/labs.

Data Ethics Research Initiative

Launched by Accenture's Technology Vision team, the Data Ethics Research Initiative brings together leading thinkers and researchers from Accenture Labs and over a dozen external organizations to explore the most pertinent issues of data ethics in the digital economy. The goal of this research initiative is to outline strategic guidelines and tactical actions businesses, government agencies, and NGOs can take to adopt ethical practices throughout their data supply chains.

About Accenture

Accenture is a leading global professional services company, providing a broad range of services and solutions in strategy, consulting, digital, technology and operations. Combining unmatched experience and specialized skills across more than 40 industries and all business functions—underpinned by the world's largest delivery network—Accenture works at the intersection of business and technology to help clients improve their performance and create sustainable value for their stakeholders. With approximately 375,000 people serving clients in more than 120 countries, Accenture drives innovation to improve the way the world works and lives. Visit us at www.accenture.com.

Learn more: www.accenture.com/DataEthics