

Software

JMP's Five Commandments of AI/ML Design

INDUSTRY OVERVIEW

This report addresses the ethics of artificial intelligence (AI) and machine learning (ML) design. In particular, we think CEOs can help drive ethical design practices in this area by sharing five principles with their product development and engineering teams. We refer to these principles as the five commandments of AI/ML design:

- Do not violate people's privacy
- Be honest
- Produce explainable results
- Be fair
- Be safe

In the body of this note, we discuss these principles in more detail.

Patrick Walravens

pwalravens@jmpsecurities.com
(415) 835-8943

Greg McDowell

gmcowell@jmpsecurities.com
(415) 835-3934

Mathew Spencer

mspencer@jmpsecurities.com
(415) 835-8930

JMP's five commandments of
AI/ML design.

DO NOT VIOLATE PEOPLE'S PRIVACY

The first commandment is “do not violate people’s privacy.” Privacy violations are perhaps the most likely ethical mistake when designing learning algorithms. It seems simple enough, but the problem is that “privacy is contextual, and situational, and not reducible to a simple public/private binary.”¹

Learning algorithms are fueled by data and most data represents or impacts people. While ethical considerations about personal data are usually characterized as issues of privacy, a broader issue is often data access. As the Institute of Electrical and Electronics Engineers (IEEE) put it, people need “to have the right to define access and provide informed consent with respect to the use of their personal data.”² The IEEE continued by saying that:

“At a minimum, consumers must have the transparency necessary to make informed decisions about whether to share their data and how it can be used. Consumers ought to have clear information, not opaque policies and complex click-through consent pages. The tech industry has an obligation to respond to widespread and growing concerns over data privacy and security and to restore the public trust. The status quo no longer works.”²

Facebook (FB, MO, \$235 PT, Josey) CEO Mark Zuckerberg learned this lesson the hard way, as his company’s partnership with Cambridge Analytica resulted in widespread privacy violations and Mr. Zuckerberg spending two days testifying in front of the Congress.

Figure 1 below shows Mr. Zuckerberg testifying before Congress.

FIGURE 1. Mark Zuckerberg Testifies Before Congress



Source: C-SPAN

Cambridge Analytica was incorporated in December 2013 in Delaware. In June 2014, entities related to Cambridge Analytica collaborated with Aleksandr Kogan, a psychologist at the University of Cambridge, to create an application for use on Facebook. The application was called “This Is Your Digital Life.” It was a personality quiz that visitors logged into using their Facebook credentials.

¹ Zook M, Barocas S, Boyd, Crawford K, Keller E, Gangadharan SP, et al. (2017) Ten simple rules for responsible big data research. PLoS Comput Biol 13(3).

² Ethically Aligned Design, Version 2 (EADv2) | IEEE Standards Association.” EADv2 - Ethics in Action.

The quiz gave users what is called an OCEAN score – referring to openness, conscientiousness, extroversion, agreeableness, and neuroticism. At the end of the quiz, about 300K users gave the This Is Your Digital Life application permission to access their Facebook profiles. Once the application had access to the profiles, it extracted the personal information of the users and their friends.³ In the end, Cambridge Analytica reportedly extracted the personal data of over 50M people.⁴

Mr. Zuckerberg acknowledged that Facebook fell short as he told Congress, “We didn’t take a broad enough view of our responsibility and that was a big mistake.”⁵

How should companies think about personal data to avoid these types of breaches? In his book, *The Master Algorithm*, Pedro Domingo, a professor of computer science at the University of Washington, provides a helpful framework for data sharing by discussing four kinds of data:

- Data to be shared with everyone;
- Data to be shared with friends or coworkers;
- Data to be shared with various companies (wittingly or not); and
- Data not to be shared at all.

The first category, data to be shared with everyone, includes things such as reviews for Yelp (YELP, MP, Josey), Amazon (AMZN, MO, \$1,840 PT, Josey), or TripAdvisor; resumes on LinkedIn (MSFT, NC); Tweets on Twitter (TWTR, MP, Josey); and blog posts.

The second category, data to be shared with friends and coworkers, includes updates and pictures shared on services such as Instagram or Facebook. Mr. Domingo commented:

“You share updates and pictures with your friends on Facebook, and they with you. But everyone shares their updates and pictures with Facebook. Lucky Facebook: it has a billion friends. Day by day it learns a lot more about the world than any one person does...The only problem is that Facebook is also free to do things with the data and the models that are not in your interest, and you have no way to stop it.”⁶

The third category, data to be shared with various companies, might include searches on Google (GOOGL, Market Outperform, \$1,235 PT, Josey), phone calls with Verizon (VZ, NC), groceries at Whole Foods, and financial transactions with Bank of America (BAC, NC).

The last type of data, according to Professor Domingo, is data not to be shared at all, such as health records. Although, in some cases, perhaps people would share this data if they could find a way to guarantee it would not be misused, or determine that sharing it would be vital for a given reason.

³ Staff, Wired. “The Cambridge Analytica Story, Explained.” Wired, Condé Nast, 22 Mar. 2018.

⁴ Greenfield, Patrick. “The Cambridge Analytica Files: the Story so Far.” The Guardian, Guardian News and Media, 25 Mar. 2018.

⁵ Timberg, Craig, and Tony Romm. “Facebook CEO Mark Zuckerberg to Capitol Hill: ‘It Was My Mistake, and I’m Sorry.’” The Washington Post, WP Company, 9 Apr. 2018.

⁶ Domingos, Pedro. *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*. Basic Books, a Member of the Perseus Books Group, 2018.

BE HONEST

The second commandment is to be honest. Honesty is important because it maintains trust in the system. Honesty, in part, means being transparent in interactions with humans. Users need to know when they are dealing with an intelligent system rather than a human. Otherwise, trust rapidly dissipates. A recent failure in transparency came from Google CEO Sundar Pichai's keynote presentation at the Google I/O developer conference on May 8. During that keynote, Mr. Pichai demonstrated the company's latest AI initiative – Google Duplex. He played a recording in which a Google Bot powered by Duplex called a hair salon and booked an appointment with a woman at the salon.

The conversation between the bot and the human went like this:

- **Human:** Hello, how can I help you?
- **Bot:** Hi, I'm calling to book a women's haircut for our client. Umm... I'm looking for something on May third.
- **Human:** Sure, give me one second.
- **Bot:** Mm-hmmm.
- **Human:** Sure, what time are you looking for, around?
- **Bot:** At 12:00 p.m.
- **Human:** We do not have a 12:00 p.m. available. The closest we have to that is a 1:15.
- **Bot:** Do you have anything between 10:00 a.m. and um... 12:00 p.m.?
- **Human:** depending on what service she would like... what service is she looking for?
- **Bot:** Just a woman's haircut, for now.
- **Human:** Okay, we have a 10 o'clock.
- **Bot:** 10am is fine.
- **Human:** Okay, what's her first name?
- **Bot:** The first name is Lisa.
- **Human:** Okay, perfect, so I will see Lisa at 10 o'clock on May third.
- **Bot:** Okay, great, thanks.
- **Human:** Great, have a great day, bye.⁷

⁷ Google I/O Keynote, May 8, 2018.

Figure 2 below shows Google CEO Sundar Pichai presenting his keynote at the Google I/O Developer Conference.

FIGURE 2. Google CEO Sundar Pichai Keynote at Google I/O Developer Conference



Source: Google website

This demonstration failed at honesty because the human being who answered the phone at the salon had no idea she was talking to a computer program. In fact, with its beautifully timed “ums” and its elongated pronunciation of certain words, the Google Bot seemed to be specifically designed to trick humans into thinking that they are actually talking to another human being.

Two days after this demonstration and after a number of journalists pointed out its unethical nature, a Google spokesperson stated, “We understand and value the discussion around Google Duplex – as we’ve said from the beginning, transparency in the technology is important.”⁸

On May 8, Google published a blog titled, “Google Duplex: An AI System for Accomplishing Real-World Tasks over the Phone,” which described some of the underlying technology behind the salon appointment call Mr. Pichai demonstrated in his keynote. One of the public comments highlighted the issue well by stating, “humans in general don’t mind talking to machines so long as they that they’re doing so.”

“We understand and value the discussion around Google Duplex – as we’ve said from the beginning, transparency in the technology is important.”

- Google spokesperson

⁸ Fagan, Kaylee. “Google’s Eerily Realistic New AI Will Identify Itself When Talking to People, Says Google.” Business Insider, 11 May 2018.

This comment is shown in Figure 3 below.

FIGURE 3. Comments on Google AI Blog



Lauren Weinstein via Google+ 1 month ago - Shared publicly

A couple of problems spring immediately to mind. First, the use of embedded 'uh's and other artifacts to try fool the listener into believing that they are speaking to a human may well engender blowback as these systems are deployed. My sense is that humans in general don't mind talking to machines *so long as they know that they're doing so*. I anticipate significant negative reactions by many persons who ultimately discover that they've been essentially conned into thinking they're talking to a human, when they actually were not. It's basic human nature -- an area where Google seems to have a continuing blind spot. Another problem of course is whether this technology will ultimately be leveraged by robocallers (criminal or not) to make all of our lives even more miserable while enriching their own coffers.

Show less

+116 · Reply

Source: Google AI Blog

PRODUCE EXPLAINABLE RESULTS

The third commandment is that intelligent systems should be able to produce explainable results.

Ryan Welsh, the CEO of Kyndi, a privately-held AI software company based in San Mateo, California, defines “explainable” as “artificial intelligence systems that can justify the reasoning behind their conclusions and results, allowing the user to have full confidence in the system's outputs.”⁹ The challenge is that many popular and effective machine learning techniques have traded off explainability for performance. Often the best performing methods, such as deep learning, are the least transparent, while the methods providing the best explanations, like decision trees, are not as accurate.¹⁰

This rule is, admittedly, perhaps not as broadly applicable as the other four. Take certain cases of image detection as an example of where explainability may be less important. For example, in 2012 Quoc V. Le and several Google colleagues working with Stanford's Andrew Ng published a seminal paper about a large scale deep learning simulation on 16,000 cores using a billion parameters that, among other things, could detect cats in videos. How important is it to explain why exactly the program detected a cat? In this example, perhaps not too important.

On the other hand, organizations in a number of areas are often unable to deploy AI if it is not explainable. Three of these areas include financial services, government, and healthcare.

In financial services, for example, credit card companies such as Capital One (COF, NC) would like to use machine learning to decide who is granted a credit card “but it cannot do that because the law requires companies to explain the reason for any such decision to a prospective customer.”¹¹

⁹ Welsh, Ryan. “What Is Explainable AI™? – Kyndi.” Kyndi, 26 Apr. 2018.

¹⁰ Holzinger, Andreas et al. “What Do We Need to Build Explainable AI systems for the Medical Domain?”

¹¹ Knight, Will. “The Financial World Wants to Open AI's Black Boxes” 13 Apr. 2017.

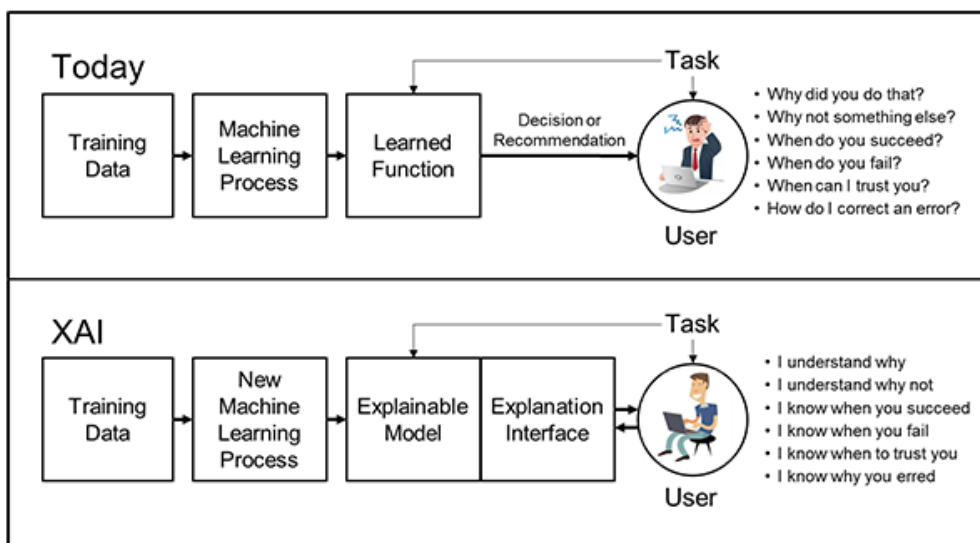
Adam Wenchel the VP of AI and Data Innovation at Capital One has commented that “we’ve got a whole work stream of experts working on explainable AI and the ability to articulate how decisions are made.”¹²

Government is another area that often requires explainable AI. Max Tegmark gives the example of robo-judges, which in theory could be programmed to adjudicate certain types of cases. But as he points out, “If defendants wish to know why there were convicted, shouldn’t they have the right to a better answer than “we trained the system on lots of data, and this is what it decided?”¹³

Besides the judiciary, another part of government that is very interested in explainable AI is the Department of Defense. DARPA, the Defense Advanced Research Projects Agency, has commented that the effectiveness of AI systems is limited by their inability to explain their actions and decisions to human users. The issue here of course is autonomous weapons systems: “Explainable AI—especially explainable machine learning—will be essential if future warfighters are to understand, appropriately trust, and effectively manage an emerging generation of artificially intelligent machine partners.”¹⁴

Figure 4 below shows the idea behind DARPA’s Explainable Artificial Intelligence (XAI) initiative.

FIGURE 4. DARPA’s XAI Initiative



Source: www.darpa.mil

A third example is medicine. Patients can benefit from having the clinical decision making skills of medical professionals complemented by sophisticated machine learning systems. A system might even overrule the medical professional in the interest of patient safety. Medical professional can also see benefits in education, research, and decision making. However, patients and doctors will still need to understand why a system suggests a certain course of treatment or offers a particularly diagnosis. “The human expert must still have a chance, on demand, to understand and retrace the machine decision process.”¹⁵

¹² New, Joshua. “5 Q’s for Adam Wenchel, Vice President of AI and Data Innovation at Capital One” Center for Data Innovation, 15 Dec. 2017.

¹³ Page 106 “LIFE 3.0: Being Human in the Age of Artificial Intelligence, by Max Tegmark, Alfred A. Knopf, 2017.

¹⁴ <https://www.darpa.mil/program/explainable-artificial-intelligence>.

¹⁵ Holzinger, Andreas et al. “What Do We Need to Build Explainable AI systems for the Medical Domain?”

BE FAIR

The fourth commandment is to be fair. CEOs should remind their teams that learning algorithms, and the results from those should be fair and free of biases. Biased results from algorithms, regardless of the programmer's intentions, can lead to unintended discrimination or perpetuate pre-existing discriminatory behaviors, particularly in areas such as recruiting, hiring, and promotions. Conversely, when used correctly, these kinds of algorithms can be a powerful tool to minimize bias in the workplace and elsewhere.

While AI systems are often intended to streamline processes in an unbiased way, there are a number of examples where based on the underlying data, AI actually reinforces stereotypes or misrepresents reality, sometimes in offensive ways. Examples include:

- Certain word association algorithms being more likely to associate the word, "honorable" with men;¹⁶
- Passport approval software in New Zealand misidentifying Asian applicants, asking them to retake their photos because their eyes were not open;¹⁷
- The case of Jacky Alcine, a software developer, who tweeted a photo of Google's recognition software labeling her and her African-American friends as Gorillas;¹⁸ and
- Researchers at University of Virginia, who confirmed that photo recognition software correctly identifies a spatula when it is on a surface, but when it is being held by a human, it is more likely to label them as female, regardless of their actual sex.¹⁹

Often these biased results come not from the algorithms itself, but from the training data. For example, the University of Virginia researchers found that the underlying datasets for certain tasks, such as cooking, contained significant gender bias, and that some algorithms actually amplify the bias.

While these stories above showcase the shortcomings of AI in relation to bias, on a more positive note, AI can also be used to offset it. The researchers at the University of Virginia suggest AI/ML system designers identify, evaluate, and quantify the amount of bias in their systems and then implement additional algorithms to correct the bias introduced into the system, including those from the underlying dataset.

An area in which algorithmic fairness is particularly important is talent management – that is, recruiting, hiring, firing, training and promoting employees. One company that seeks to improve talent management processes using AI is Eightfold.ai, a privately held AI software company based in Mountain View, California. Eightfold CEO, Ashutosh Garg, has commented that, "If algorithms and models are trained on past actions – we could end up perpetuating the prior bias. Eightfold specifically uncovers bias and low diversity – assuming best intentions – and then provides the ability to do the course corrections."

"If algorithms and models are trained on past actions - we could end up perpetuating the prior bias. Eightfold specifically uncovers bias and low diversity - assuming best intentions - and then provides the ability to do the course corrections."

- Eightfold CEO, Ashutosh Garg

¹⁶ "Machine-Learning Algorithm Measures Changes in Gender, Ethnic Bias." Stanford Bio-X, Stanford.

¹⁷ Griffiths, James. "New Zealand Passport Robot Thinks This Asian Man's Eyes Are Closed." CNN, Cable News Network, 9 Dec. 2016.

¹⁸ Simonte, Tom. "When It Comes To Gorillas, Google Photos Remain Blind." Wired, 11 Jan. 2018.

¹⁹ Zhao, Jieyu et al. "Reducing Gender Bias Amplification Using Corpus-Level Constraints." University of Virginia.

BE SAFE

The final commandment is to be safe. Safety in the context of learning algorithms and AI refers to “unintended and harmful behavior that may emerge from poor design of real-world AI systems.”²⁰ In his book *Life 3.0: Being Human in the Age of Artificial Intelligence*, Max Tegmark described four main areas of AI safety:

- Verification;
- Validation;
- Control; and
- Security.

These are the same four problem areas that have traditionally affected software systems, but AI exacerbates these problems because the systems keep learning and changing their behavior.

Below, we discuss each area and provide some examples. While these examples are older and deal with traditional software, we think the lessons are highly relevant to today’s AI/ML system design.

Verification

Verification means proving the correctness of the underlying algorithm. It means “ensuring that the software fully satisfies the expected requirements.” Or as Tegmark neatly summarizes, verification asks the question: “Did I build the system right?”²¹

One example of verification that Mr. Tegmark brings up is Knight Capital, which lost about \$460M on August 1, 2012 in a 45-minute period, due to a verification error in its automated routing system for equity orders, called SMARS. Among other things, SMARS received “parent” orders internally and then based on available liquidity would send “child” orders to external venues for execution. Knight inserted new code into its SMARS to enable its customer participation in the Retail Liquidity Program at the New York Stock Exchange.

During a 45-minute period following the installation, SMARS sent millions of orders resulting in 4 million transactions in 154 stocks. These transactions resulted in a \$3.5B net long position in 80 stocks, a \$3.2B net short position in 74 stocks and a realized loss of \$460M. Ultimately, Knight’s loss from these unwanted positions, forcing the company to sell off a majority of its share in order to avoid bankruptcy.²²

The SEC issued an Order on October 16, 2013 (Commission, 2013) that concluded:

“Knight did not have technology governance controls and supervisory procedures sufficient to ensure the orderly deployment of new code or to prevent the activation of code no longer intended for use in Knight’s current operations, but left on its servers that were accessing the market.”²³

²⁰ Page 94 “ LIFE 3.0: Being Human in the Age of Artificial Intelligence, by Max Tegmark, Alfred A. Knopf, 2017.

²¹ Page 96.” LIFE 3.0: Being Human in the Age of Artificial Intelligence, by Max Tegmark, Alfred A. Knopf, 2017.

²² Knightmare: A DevOps Cautionary Tale.” Doug Seven, 17 Apr. 2014.

²³ In the Matter of Knight Capital LLC.” Sec.gov, 16 Oct. 2013.

Validation

In contrast to verification, validation asks if the specifications of a system are appropriate for the task it was built to accomplish. In other words, validation asks, “did I build the right system?” Tegmark provides an example of incorrect validation in the case of the “Flash Crash” on May 6, 2010. On this day, the Dow Jones Industrial Average plunged over 600 points in a few minutes, only to rebound approximately 40 minutes later and regain most of its losses. During this period, share prices for some companies became irrational, varying between \$0.01 and \$100,000. According to Tegmark, neither software nor computer malfunctions created the problem. Instead, the systems which generate automatic algorithmic trades were operating in unexpected situations where certain assumptions were not valid, such as a securities’ price. According to Tegmark, crashes such as these would not have been avoided by simply checking over the systems in place for errors.

According to the SEC’s report on the crash, the catalyst occurred when a large mutual fund initiated a sizeable trade through a sell algorithm. That sell algorithm was based on the trading volume across the previous minute; it did not take into account price or time. In addition, this trade was placed in a market that was already under stress, as the Dow had traded down 2.5% before the trade was placed. High-frequency traders initially absorbed the selling pressure. These traders bought and subsequently sold shares, which increased volume and accelerated the pace of the mutual fund’s sell algorithm based on volume. In less than five minutes, the S&P 500 had declined 6%.²⁴ Some traders were able to note the irregularities and halt their trading, and some trading systems had automatic trading pauses built into their systems based on the market events. Each of these factors contributed to decreased liquidity and exacerbated the decline. In response to abnormally large price declines, out-of-balance positions, and a lack of liquidity, some market makers began to generate automatic quotes (referred to as “stub quotes”) at levels far away from the current market. The market makers did this to protect themselves while, at the same time, satisfying continuous two-sided quoting obligations. Other market participants, including investors, had their trades executed in the disrupted markets, some at prices that were based on quotes that were never intended to reflect the fundamental value of the corresponding securities.

As this example illustrates, validation goes beyond simply checking for errors in the code; it deals with the design of a system itself.

Control

Control is the idea that humans need to be able to interact with a technology or system effectively and accurately. This is especially important for so-called Human-In-The-Loop systems. The ability to monitor and alter such systems is paramount to the system’s practicality.

Monitoring and altering can take many forms. In the case of British car ferry Herald of Free Enterprise, it can be as simple as a warning light on the captain’s dashboard to let him know that the ship’s bow doors are open. In that 1987 incident, the captain failed to see that the bow doors were not secured, and there was no mechanism on the bridge in place to alert him to that fact. This tragedy resulted in the boat capsizing and 193 people dying on board.²⁵ The simple addition of a warning light could have saved these lives.

²⁴ “Findings Regarding The Market Events Of May 6, 2010.” Securities and Exchange Commission, 30 Sept. 2010.

²⁵ “MS Herald of Free Enterprise.” Ship Disasters.

Figure 5 below shows the MS Herald of Free Enterprise.

FIGURE 5. MS Herald of Free Enterprise



Source: Wikipedia

Similarly tragic, when a plane crashed into the Vosges Mountains in France in 1992, it was because the graphical user interface on the autopilot did not display which mode it was in. Therefore, when the pilot entered “33,” trying to set a descent angle of 3.3 degrees, the autopilot read it as setting the speed to 3,300 feet per minute, or 37.5 miles per hour.²⁶ Had there simply been a larger screen that displayed the mode the autopilot was set to, the tragedy may have been avoided.

The loss of life in both of these examples pales in comparison to what came close to occurring in October of 1962. During the Cuban Missile Crisis, Soviet submarine B-59 was in a standoff with 12 U.S. Navy ships off the coast of Florida. The submarine had lost contact with Moscow days earlier, and when the American ships began dropping weak depth charges in an attempt to get the sub to surface and leave, the Russian crew thought that they were under attack and war had broken out. The submarine was authorized to launch nuclear-armed torpedoes in this event, which required the authorization of the three commanding officers. Two officers signed off to launch, but one man, Vasili Arkhipov, did not.²⁷ Nuclear war may have been averted by one single man refusing to go along with the rest of his fellow officers. Had this been an AI-controlled ship with no human interference, World War III may well have begun, as an AI-controlled ship risks not having the capabilities to go against protocols and not fire the torpedo.

Mr. Arkhipov received the first "Future of Life Award," which was presented posthumously to his family in 2017. This award, from the Future of Life Institute, recognizes exceptional measures, often performed despite personal risk and without obvious reward, to safeguard the collective future of humanity. The Future of Life Institute was founded in 2014 by Max Tegmark and four others.

Figure 6 below shows Soviet submarine B-59.

²⁶ Phillips, Don. "Manslaughter Trial of 6 Starts in 1992 Air Crash." The New York Times, The New York Times, 2 May 2006.

²⁷ Davis, Nicola. "Soviet Submarine Officer Who Averted Nuclear War Honoured with Prize." The Guardian, Guardian News and Media, 27 Oct. 2017.

FIGURE 6. Soviet Submarine B-59

Source: U.S. Naval Archives

Even if a system is verified and validated correctly, humans need to be able to work within the system to make sure it is accomplishing its function. Poor design, poor implementation, or no human interaction at all, as seen in the examples discussed above, are hallmarks of poor control.

Security

Security is to protect against attacks to software systems and data, including from malware and hacks. While the other AI risks outlined in this section relate to unintentional mistakes, this risk is intentional.

One example of a security breach that Mr. Tegmark brings up is what became known as the ILOVEYOU computer worm. On May 5, 2000, people began receiving emails with the subject line "ILOVEYOU" and an attachment titled, "LOVE-LETTER-FOR-YOU.txt.vbs". Users who clicked on the attachment unwittingly launched a script that damaged their computers and resent the email to everyone in their address book. The malware eventually infected about 50M computers and caused over \$5B in damages.²⁸

Another example came in September 2016, when Yahoo disclosed it had been subject to the largest data breach in history. The breach initially compromised the personal information of 500M users according to Yahoo. It later revised its estimate to over 1B user accounts. In October of 2017, Yahoo again revised that estimate, saying that, in fact, all 3B user accounts had been compromised.²⁹

There are numerous other examples of major data breaches in the last decade. In 2013, Target (TGT, NC) had almost 110M of its customer's credit card numbers, in addition to more personal information, compromised. Equifax was hacked in 2017, allowing access to upwards of almost 146M social security numbers. In 2014, hackers used stolen eBay (EBAY, MP, Josey) employee credentials to access eBay's user database, compromising the information of 145M users on the site.³⁰

As AI's role in society increases, so too does the need for greater security. While AI can be used to improve security it can also be used to find new vulnerabilities and more sophisticated attacks.

In conclusion, while AI/ML design is a complicated and deeply technical area, we believe CEOs can help drive ethical design in AI/ML by sharing these principles with their teams early and often.

²⁸ Tegmark, Max. LIFE 3.0: Being Human in the Age of Artificial Intelligence. Alfred A. Knopf, 2017.

²⁹ Yahoo. "An Important Message to Yahoo Users on Security." Edgar, Securities and Exchange Commission.

³⁰ Armerding, Taylor. "The 17 Biggest Data Breaches of the 21st Century." CSO Online, InfoWorld, 26 Jan. 2018.

JMP FACTS AND DISCLOSURES

Analyst Certification:

The research analyst(s) who prepared this report does/do hereby certify that the views presented in this report are in accordance with my/our personal views on the securities and issuers discussed in this report. As mandated by SEC Regulation AC no part of my/our compensation was, is or will be directly or indirectly related to the specific views or recommendations expressed herein. This certification is made under the obligations set forth in SEC Regulation AC. Any other person or entity may not use it for any other purpose. This certification is made based on my/our analysis on the date of this report's publication. I/We assume no obligation to update this certification to reflect any facts, circumstances, or events that may subsequently come to my/our attention. Signed Patrick Walravens, Greg McDowell and Mathew Spencer

JMP Securities Disclosures:

JMP Securities currently makes a market in the securities of Amazon.com, Inc., eBay Inc., Facebook, Inc., Alphabet Inc., Twitter, Inc. and Yelp Inc.

JMP Securities expects to receive OR intends to seek compensation for investment banking services from Amazon.com, Inc., eBay Inc., Facebook, Inc., Alphabet Inc., Twitter, Inc. and Yelp Inc. in the next 3 months.

JMP Securities Investment Opinion Definitions:

Market Outperform (MO): JMP Securities expects the stock price to outperform relevant market indices over the next 12 months.

Market Perform (MP): JMP Securities expects the stock price to perform in line with relevant market indices over the next 12 months.

Market Underperform (MU): JMP Securities expects the stock price to underperform relevant market indices over the next 12 months.

JMP Securities Research Ratings and Investment Banking Services: (as of June 25, 2018)

JMP Rating	Regulatory Equivalent	# Co's Under Coverage	% of Total	Regulatory Equivalent	# Co's Under Coverage	% of Total	# Co's Receiving IB Services in Past 12 Months	
							# Co's Receiving IB Services in Past 12 Months	% of Co's With This Rating
MARKET OUTPERFORM	Buy	270	62.21%	Buy	270	62.21%	69	25.56%
MARKET PERFORM	Hold	156	35.94%	Hold	156	35.94%	26	16.67%
MARKET UNDERPERFORM	Sell	6	1.38%	Sell	6	1.38%	0	0%
COVERAGE IN TRANSITION		1	0.23%		1	0.23%	0	0%
RATING SUSPENDED		0	0.00%		0	0.00%	0	0%
TOTAL:		434	100%		434	100%	95	21.89%

Stock Price Chart of Rating and Target Price Changes:

Note: First annotation denotes initiation of coverage or 3 years, whichever is shorter. If no target price is listed, then the target price is N/A. In accordance with FINRA Rule 2241, the chart(s) below reflect(s) price range and any changes to the rating or price target as of the end of the most recent calendar quarter. The action reflected in this note is not annotated in the stock price chart. Source: JMP Securities.

In order to obtain these (6 or more) stock price charts or additional applicable disclosures and information concerning JMP's recommendations of companies under coverage mentioned in this report, please contact JMP Securities at (877) 263-1333 or visit www.jmpsecurities.com

JMP Disclaimer:

JMP Securities LLC (the "Firm") compensates research analysts, like other Firm employees, based on the Firm's profitability, which includes revenues from the Firm's institutional sales, trading, and investment banking departments as well as on the quality of the services and activities performed that are intended to benefit the Firm's institutional clients. These data have been prepared by JMP Securities LLC for informational purposes only and are based on information available to the public from sources that we believe to be reliable, but we do not guarantee their accuracy or completeness. Any opinions and projections expressed herein reflect our judgment at this date and are subject to change without notice. These data are neither intended nor should be considered as an offer to sell or a solicitation or a basis for any contract for the purchase of any security or other financial product. JMP Securities LLC, its affiliates, JMP Group LLC, Harvest Capital Strategies LLC, and their respective partners, directors, officers, and associates may have a long or short position in, may act as a market maker for, or may purchase or sell a position in the securities mentioned herein. JMP Securities LLC or its affiliates may be performing, have performed, or seek to perform investment banking, advisory, or other services and may have acted as manager or co-manager for a public offering of securities for any company mentioned herein. The reader should assume that JMP Securities LLC will solicit business from the company covered in this report. Members of our Sales and Trading Department provide oral and/or written market opinions and trading strategies to our clients that reflect their personal opinions about stocks that are the subject of the firm's research reports. Our research analysts discuss trading strategies with clients that sometimes reflect short-term expectations for the price of the securities that are the subject of research reports. These trading strategies are distinct from the analysts' fundamental rating for the stock, which is based upon the analysts' view compared to other stocks under coverage for the relevant time period. © Copyright 2018. All rights reserved by JMP Securities LLC. JMP Securities LLC is a member of FINRA, NASDAQ, and SIPC.

Jeffrey H. Spurr
 Director of Research
 (415) 835-3903

RESEARCH PROFESSIONALS

FINANCIAL SERVICES

Alternative Asset Managers

Devin Ryan (212) 906-3578
 Brian McKenna (212) 906-3545

Commercial & Specialty Finance

Christopher York (415) 835-8965
 Thomas Wenk (415) 835-8962

Consumer Finance

David M. Scharf (415) 835-8942
 Jeff Zhang, CFA (415) 835-8948

Financial Processing & Outsourcing

David M. Scharf (415) 835-8942
 Jeff Zhang, CFA (415) 835-8948

Insurance

Matthew J. Carletti (312) 768-1784
 Karol Chmiel (312) 768-1786

Investment Banks & Brokers

Devin Ryan (212) 906-3578
 Brian McKenna (212) 906-3545

Mortgage Operating Companies

REITs: Agency, Hybrid, & Commercial Mortgage

Steven C. DeLaney (212) 906-3517
 Chris Muller (212) 906-3559
 Trevor Cranston, CFA (415) 869-4431
 Mikhail Goberman (212) 906-3543

HEALTHCARE

Biotechnology

Liisa A. Bayko (312) 768-1785
 Jonathan Wolleben (312) 768-1788
 Jason N. Butler, PhD (212) 906-3505
 Roy Buchanan (212) 906-3509
 Konstantinos Aprilakis, MD (212) 906-3503
 Simon Gruber (212) 906-3538

Healthcare Services & Facilities

Peter L. Martin, CFA (415) 835-8904

Medical Devices & Supplies

David Turkaly (212) 906-3563
 Daniel W. Stauder (212) 906-3535

Specialty Pharmaceuticals

Donald Ellis (212) 906-3507
 Nazibur Rahman (212) 906-3519

REAL ESTATE

Housing & Land Development

Peter L. Martin, CFA (415) 835-8904
 Aaron Hecht (415) 835-3963
 Doug Hansen (415) 835-8934

Property Services

Mitch Germain (212) 906-3546
 Corey DeVito (212) 906-3525

Residential Services

Aaron Hecht (415) 835-3963

REITs: Healthcare & Specialty

Peter L. Martin, CFA (415) 835-8904

REITs: Diversified, Industrial, Office, & Retail

Mitch Germain (212) 906-3546
 Corey DeVito (212) 906-3525

REITs: Residential

Aaron Hecht (415) 835-3963

TECHNOLOGY

Energy Technology & Services

Industrial Technology

Joseph Osha (415) 835-8998
 Hilary Cauley (415) 835-8996

Cybersecurity, Data Management & Comm Infrastructure

Erik Suppiger (415) 835-3918
 Michael Berg (415) 835-3914

Internet & Digital Media

Ronald V. Josey III (212) 906-3528
 Mike McGovern (415) 835-8916

Software

Patrick Walravens (415) 835-8943
 Mathew Spencer (415) 835-8930
 Greg McDowell (415) 835-3934

ADDITIONAL CONTACTS

Thomas R. Wright
 Director of Equities
 (212) 906-3599

Charles Sweeney
 Director of Sales & Trading
 (212) 906-3573

600 Montgomery Street, Suite 1100
 San Francisco, CA 94111
www.jmpsecurities.com