GETTY IMAGES

TABLE *of* EXPERTS

The Impact of AI

Bay Area experts discuss the meaning of Artificial Intelligence and how it will affect the economy and society

At a Dec. 9 forum hosted by the San Francisco Business Times, panelists Christian Fjeld, vice president at ML Strategies, James Meadow, AI and Machine Learning Principal at Slalom, and Tom Mescall, partner and head of consulting at Armanino, discussed how businesses can leverage AI to drive innovation, the impact AI

MEET THE PANELISTS
•
NEXT PAGE

will have on the economy and on society, and the meaning of artificial intelligence itself. The discussion was moderated by Mary Huss, Bay Area president and publisher of the San Francisco Business Times and Silicon Valley Business Journal. The panelists' comments have been edited for length and for clarity.

Q: How should business leaders think about artificial intelligence?

James Meadow, AI and Machine Learning Principal at Slalom: Artificial intelligence is a piece of software that can make a complex, human-like decision. First, you need to distinguish between machine learning and AI. Machine learning is a technical exercise. A statistician or a data scientist or an engineer trains a model to answer a question – that's machine learning. It's about finding ways to make human-like decisions using a computer algorithm.

AI, on the other hand, is a more complex engineered process. AI is also more cross-disciplinary than ML. It requires design; it requires infrastructure and it requires very clean, predictable data to pull it off correctly. Machine learning is a component of AI. But when you talk to your phone, you're talking to an AI. You're interacting with a heavily engineered system with many different parts. At the very bottom of that system is a machine learning algorithm that listens to what you say, parses what you say, and then another machine learning algorithm that decides what to do with that information.

Christian Fjeld, vice president at ML Strategies: Speaking from a policy perspective, the Congressional Research Service (a component of the Library of Congress that produces research and policy papers for members of Congress and their staff) defines AI as "computerized systems that work and react in ways commonly thought to require intelligence, such as solving complex problems in real-world situations." On the other hand, the Association of Artificial Intelligence states that AI researchers are seeking to understand "the mechanisms underlying thought and intelligent behavior and their embodiment in machines." While both are good definitions, Congress and policy-makers tend to focus on real-world situations.

Tom Mescall, partner and head of consulting at Armanino: Looking at it from a business perspective, you can boil AI down to three technological pillars. The first is RPA, or robotic process automation. That's when a bot is doing something

Continued on next page

Sponsored by:

armanino   **ML** STRATEGIES   **M** MINTZ   slalom

TABLE *of* EXPERTS

PANEL DISCUSSION:

The Impact of AI

Bay Area experts discuss the meaning of Artificial Intelligence and how it will affect the economy and society

that was previously done by a human. It could be onboarding a new employee, paying a bill, or adding a new vendor to the system. It's any situation where in the past humans were doing the task, but now a bot takes care of it. It's not about replacing employees necessarily; it's about gaining efficiencies.

The second component is virtual assistants. The key technology behind virtual assistants is natural language processing. There's been a lot academic discussion around what really counts as natural language processing, but what business leaders need to know is how it manifests itself in business. That's when a human can interact with a computer and get back information. Alexa and Siri are some obvious consumer examples, but there's also lot of business applications. A customer service center, for instance, could involve people on your website interacting with virtual assistants.

Finally, there's predictive analytics, which is arguably the term most associated with AI. "Machine learning," "deep learning" and "big data," are all related terms you'll hear associated with predictive analytics. What they're all getting at is the idea of gleaning something novel from the data. An example of that might be customer churn: Is a customer going to leave my business any time soon? Predictive analytics asks, "What are the characteristics of the data that would inform me about what's going to happen in the future?"

Obviously, when you talk to scientists and IT professionals, the topic of AI becomes a lot more technical. And outside of a business context, AI is something else – it applies to efforts such as autonomous vehicles, curing cancer, mapping the human genome, etc. But narrowing it down and framing it for a business person, AI can be understood as being composed of those three pillars.

Meadow: It's also important to distinguish between generalized AI and specialized AI. Generalized AI, often referred to as artificial general intelligence, is what the public tends to think of as AI: a robot that can move and act in human-like ways and learn things on-the-fly. But most businesses are interested in specialized AI, which refers to algorithms and processes that can complete very specific tasks the way a human used to do them, but far more efficiently.

Q: What kind of interactions is AI currently capable of, and where is it going?

Mescall: To enable the future of customer interactions, natural language processing modules will need to have great image recognition capabilities, text recognition, and context recognition. Because it's not just the text that matters, it's the context in which it's said. And that's where AI is headed: predicting emotions. We are already seeing how businesses are interacting with artificial intelligence in our AI Lab where members are exploring how to automate expense reports, improve supply chain planning and predict customer churn rate. The AI Lab allows customers to evaluate an AI proof of concept fairly quickly to allow for further experimentation with low risk. The exploration is en-

---

MODERATOR

MARY HUSS

BAY AREA PRESIDENT AND PUBLISHER, SAN FRANCISCO BUSINESS TIMES AND SILICON VALLEY BUSINESS JOURNAL

Mary Huss is Bay Area president and publisher of the San Francisco Business Times, the No. 1 source for local business news and information in the San Francisco Bay Area, as well as its sister paper, Silicon Valley Business Journal. In her tenure, the paper has grown to be one of the largest and most profitable publications in the American City Business Journals portfolio. Huss first joined American City Business Journals in 1980 as director of circulation for the St. Louis Business Journal. She then became publisher of the Philadelphia Business Journal before establishing her current post at the Business Times. Outside of her publishing duties, she serves on the boards for the Bay Area Council, the San Francisco Chamber of Commerce and the San Francisco State University Foundation.

MEET THE PANELISTS

TOM MESCALL

PARTNER-IN-CHARGE OF CONSULTING, ARMANINO

Tom Mescall is partner-in-charge of consulting at Armanino, one of the top 25 largest independent accounting and consulting firms in the country. He is also the founding force behind the firm's AI Lab, which was just announced earlier this year. Since joining Armanino in 2004, he has grown the firm's once-small consulting group into the award-winning, high-growth practice area it is today. With more than 30 years of experience in management consulting, software development and public accounting, Mescall has worked extensively with ERP, eCommerce, business intelligence, artificial intelligence and CRM systems. He is a sought-after speaker for programs focused on these technologies, their future trends, and their impact on strategic planning for growth-oriented companies.

JAMES MEADOW

AI AND MACHINE LEARNING PRINCIPAL, SLALOM

James Meadow is data and analytics leader at Slalom, a consulting firm focused on strategy, technology and business transformation. Headquartered in Seattle, Washington, Slalom employees more than 7,000 employees in 30 offices across North America, Europe and the Asia Pacific region. Meadow oversees the firm's machine learning and artificial intelligence teams across Northern California, working with a wide range of technology clients. He holds a Ph.D. in genomics and machine learning from Montana State University, where his dissertation work focused on using multivariate statistical methods and machine learning techniques to distill the underlying structure of complex microbial communities. He has since has since developed and applied machine learning technology to solve problems across the forensics, health care, life sciences, marketing and public sectors.

CHRISTIAN FJELD

VICE PRESIDENT, ML STRATEGIES

Christian Fjeld is a vice president at ML Strategies, a government relations firm and wholly-owned subsidiary of Boston-based corporate law firm Mintz Levin. Based out of the firm's Washington, D.C. office, Fjeld assists a variety of business clients in their interactions with the federal government. He joined ML Strategies earlier this year. Prior to joining ML Strategies, he spent nearly 10 years in staff leadership roles with the U.S. Senate's Committee on Commerce, Science, and Transportation, serving Democratic Sens. Jay Rockefeller and Bill Nelson, as well as the current ranking member, Sen. Maria Cantwell, D-Wash. Earlier, Fjeld held policy coordinator, senior staff, and senior advisor roles with the U.S. House of Representatives' Committee on Energy and Commerce, where he worked on matters relating to the FTC, the National Highway Traffic Safety Administration, the U.S. Consumer Product Safety Commission, the Department of Commerce, and the EPA.

The Impact of AI
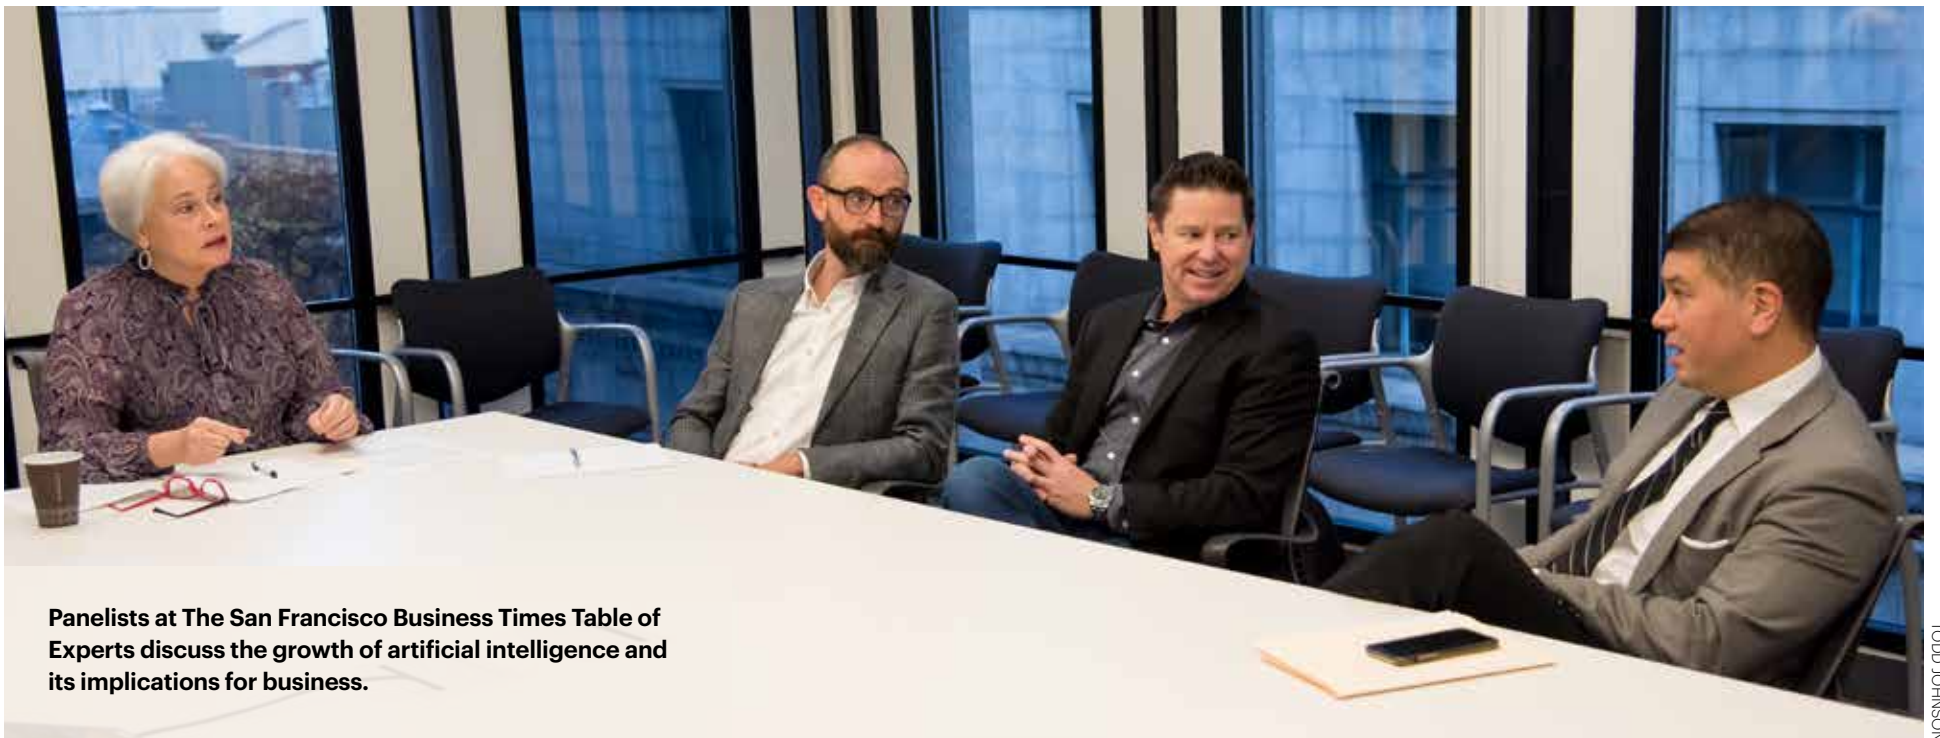
Panelists at The San Francisco Business Times Table of Experts discuss the growth of artificial intelligence and its implications for business.

abling business leaders to think more strategically on how AI can drive their strategic business processes.

Meadow: Obviously, bot interactions today can be very bumpy. But they're improving rapidly. As those interactions get better, AI will be able to take over an even greater proportion of tasks, even to the point of being an improvement over the friction that can arise when customers and customer service representatives talk past each other during a service call or chat.

Q: Is AI a threat to human workers?

Fjeld: A major public policy issue that we must grapple with, with regard to AI, are the effects it will have on the workforce. What is going to be the net effect on low-skilled labor versus high-skilled labor? Are the benefits going to be distributed equitably, or will AI benefit one class of citizens over another? Those questions are among the top on policymakers' lists when it comes to the practical impact of AI on our society.

Meadow: The prospect of better and better AI is exciting, and has the potential to change the world in positive ways. But it can also be frightening to think about from a societal perspective. For instance, a huge share of our economy consists of people driving delivery trucks. AI will potentially change that over the next decade, and that makes good sense from an efficiency and public safety perspective. However, that's a lot of people who could potentially be out of work in a few years without a secure fallback. It's hard to know how Congress is going to react to that. We're seeing ideas like the UBI (universal basic income) emerge as a response to the far-reaching effects

this is going to have. But we have a lot of work to do to refocus workers and to fundamentally rethink how our economy operates.

Mescall: In every past business revolution, the same fear has presented itself: essentially, that it's going to displace a bunch of people. But what we've seen each time is that the revolution creates more jobs, not less. During the past revolution – the advent of personal computers – we saw the same concerns. A lot of people said "computers are going to replace everybody's jobs." In fact, it created the tech industry, which today is the largest industry. It dwarfs transportation, energy and everything else. So what I predict we'll see is the economy start shifting to creating more AI-oriented jobs than we've seen up to now.

Fjeld: Washington is acutely aware that AI is the future and that we need to remain competitive in the global marketplace, particularly with China, on research and development. The reason the federal government is spending almost a billion dollars on non-defense R&D research – and likely billions more if you include classified research conducted by the Department of Defense – is that Washington is worried, understandably, that China is investing more money into artificial intelligence R&D than we are. Obviously, we don't want to lose that competitive edge. The United States has always been seen as leaders in this area. But there's a bit of a split personality of sorts, in which on the one hand, we want to nurture this field to make American businesses as competitive as possible on the global stage, but at the same time, we are a little wary about the potential for certain negative outcomes.

Q: How can AI help drive business innovation?

Meadow: When you start thinking about applying AI, you can think of three buckets: automate, amplify, and simplify. You might want to automate rote tasks – that is, tasks that are repetitive, and where people have to just sit there and do the same

thing all day long. You might want to amplify the efforts of just a few workers so that a little bit of work goes a long way. Or your goal might be to simplify a lot of the work that employees are doing. In many cases, AI can simply eliminate tasks out of an employee's work process so that

TABLE *of* EXPERTS

they can focus on activities that actually create value. If you look at your business needs through that lens, you can probably start to identify a lot of use cases. And even if it doesn't seem like machine learning or AI can tackle that problem right now, there's probably a solution that can take some of those tasks away.

AI exists on just one end of a spectrum of solutions that can be utilized to solve these problems. For instance, you might be able to write a clever query of your data, or a simple analytics solution to accomplish what might look, at first, like a machine learning problem. Or you might find a way to tap into the data on the back end of the software and automate tasks behind the scenes When those solutions become too complex or brittle, you should probably try to identify bottlenecks to your progress that can be solved with well-engineered AI.

Mescall: From a business perspective, you have to think about what common business challenges keep executives up at night. You need to go across the business landscape and think about what you can do to be a better company. When you

TODD JOHNSON

> "Machine learning," "deep learning" and "big data," are all related terms you'll hear associated with predictive analytics. What they're all getting at is the idea of gleaning something deeply relevant from the data."
>
> **Tom Mescall,** Partner-in-Charge of Consulting, Armanino

look at a company and how it's performing, what you're looking at is the sum of all the decisions that they've made previously: all of the decisions about who they've hired, all of the technology that they've deployed, etc. So it follows that if a company wants their results to be something different, then they've

got to do something different than what they've done in the past – because what they've done is what's given them their current results.

In that context, AI should be analyzed from the perspective of how it can help generate better results. It could be something as basic as noticing that you have a lot of manual processes that today's AI solutions can automate away. But it could also be about asking, "How do we engage our customers better? What information do we want about our customers that we've never had access to before, but now we can aggregate thanks to big data?" Starting with a roadmap and truly determining what you are trying to accomplish as an organization is critical in understanding key milestones and a cadence of investment to get early wins and plan for the future.

Q: What legal challenges should companies know about that arise from AI?

Meadow: The legal landscape is still catching up to the business landscape. But more than that, the research landscape is far ahead of where most of us even in the business landscape are. What that means is that we're in a position where not only do businesses have to (and should) do what's legally right, but everyone working in the field – whether you're designing algorithms, working with data, designing the front-end interaction, or some other step in the process – has to really focus on doing what's right.

To a business, that means making sure you do right by your customers. There have been some high-profile cases recently where the technology got ahead of what was considered right by customers.

They got there because designing the tech was fun. And it was really valuable, from a business perspective. But they neglected to take that necessary step back of making sure that they were doing the right thing according to their customers – to the people they're serving and whose data they're handling. So everyone in the field needs to take a very close look at what they're doing while the legal and regulatory landscape continues to lag behind where the leading edge of AI is right now. Finally, seek a contrasting opinion early and often throughout the development process when you're working on any technology with potential ethical implications

Fjeld: The regulatory landscape hasn't even caught up with the internet yet, much less AI. The big question people in Washington are asking is, does our existing regulatory architecture need to be significantly revised in order to meet the challenges of artificial intelligence, or can we use existing laws in order to protect American consumers? For instance, it is illegal to discriminate against classes of people based on their race, religion, gender or disability status. But how would the law handle an intelligent algorithm that is effectively a black box, and masks decisions that would otherwise be against the law? So in response, Democratic Senators Ron Wyden (D-OR) and Cory Booker (D-NJ) have introduced the Algorithms Accountability Act, which would require the FTC to promulgate rules to require businesses to do annual audits on algorithmic decision-making in order to ensure that there's no bias or discrimination taking place. A recently released Senate privacy draft bill has a similar provision.

There are also ethical issues that

The Impact of AI

raise concerns. For instance, Congress is currently considering autonomous vehicle legislation. That legislation could address issues such as, "Who must a car protect in case of an unavoidable accident, the passenger in the car or the pedestrian?" Right now, I think most of us just don't know the answer to that. For that reason, I expect there to be a great deal of uncertainty surrounding liability associated with the development and deployment of fully autonomous vehicles.

Similarly, the impact of AI on information and cybersecurity will likely be a double-edged sword. On the one hand, it will require fewer human eyeballs to monitor networks and to secure systems. At the same time, however, cybercriminals will be able to weaponize AI to conduct autonomous attacks on governments, businesses and critical infrastructure. There is a very uncertain regulatory landscape right now for cybersecurity when it comes to AI, and it's only going to get more complicated as technology progresses.

Mescall: All that said, I hesitate to put up any barriers to companies trying to enter the AI space. As has been frequently noted, market research shows that most business leaders believe that AI will be critical to their future; that it's going to change their industry. They believe they are going to have to do things differently in the future to succeed. But only a small fraction of companies are broadly deploying AI across their business today. So the promise is really high, but the utilization is low. What we stress to our clients is, "Don't use these legal challenges to set up road blocks as way to hamper innovation. Yes, you absolutely need to consider privacy and security and all that, but let's not let these concerns become excuses to not do anything. And additionally, there are ways for companies to get going that will minimize their risk. The real risk is staying in one place while bolder industry leaders out-innovate and, ultimately, out-compete you.

Q: What are "clean" and "dirty" data? Why is having clean data so important for AI development?

Meadow: "Dirty" data is data that is inaccurate, incomplete or inconsistent. Essentially, it's not structured in a way to be easily read and understood by a computer. For example, data could be dirty because it contains duplications. Almost all data sets contain some duplications, maybe resulting from when someone entered the same thing twice without noticing. When you train an algorithm on a data set where duplications have propagat-

TODD JOHNSON

"The regulatory landscape hasn't even caught up with the internet yet, much less AI. The big question people in Washington are asking is, does our existing regulatory architecture need to be radical revised in order to meet the challenges of artificial intelligence?"

Christian Fjeld, Vice President, ML Strategies

ed, you're essentially training it to do the wrong thing, because the data is wrong. The important thing to remember about AI is that your algorithm is only ever as good as the data used to train it. Our clients often find that cleaning the data is lot more work than they expected. In fact, a lot of what we do to prepare clients for projects is getting their data landscape clean. And that's still mainly a human process; it requires human intelligence and judgment.

There's a related concept that business leaders should also be aware of, and that's labeled data. Labeled data is a data set that describes some feature of the world, and somewhere in there is the answer to your question. For instance, if you're trying to gain insights into customer churn, in order to, say, be able to predict which of your current customers is likely to leave, you need to have a historic data set of all their relevant activities. Quite a lot of companies, however, don't have any labels on their data. That means they don't even have the answers you need to be able to train an algorithm to predict anything.

Labeling is a really big issue right now because while data is everywhere, all companies have dirty data lying around that needs attention before it can be fully utilized. What we tell clients to encourage them to clean their data

TABLE *of* EXPERTS

is that right now, data is one of the most important commodities that a business has. Spending the time and effort to clean up your data is the first step on the path to AI.

A consequence of all this is that AI is not for the faint of heart. It can be expensive, to be sure. It can also involve a lot of your work force and siphon away some of important resources. But it can ultimately pay back major dividends, if you do it right. So by all means, dip your toe in it, but know that to get serious about AI, you really need to invest time, effort and thoughtfulness into how this can help your business.

Mescal: I think every company is dealing with the problem of a non-perfect data set. Nobody has perfect data, and some data sets are obviously worse than others. This is one area where data architects and data scientists can really help: with categorization. If you've never done this before, you are not likely to succeed by yourself. You need to bring in some data architects or data scientists who have the expertise to complete this categorization. Once they've determined how your data should be categorized, less experienced employees can learn to capture the right information, and you can create work flows that reliably lead to better data sets. But really helping companies think through the categorization and how their overall landscape should look is where the expertise of a data scientist is invaluable.

Q: What role should "own" the AI function in an organization?

Mescall: Like anything else in business, there should be at least one person whose responsible for a company's AI. As a company becomes more data-oriented, they should really start looking at hiring a chief data officer.

Fjeld: It has long been federal policy, albeit somewhat voluntarily, that cyber security and data privacy should be a C-suite priority. Ideally, there should be one person who's accountable for and focuses on those issues. With CCPA (the California Consumer Privacy Act) coming into effect, that responsibility is going to be even further elevated. It will be imperative that that official interact with the chief AI officer in a company, and companies should spend some time figuring out just how those two roles will interact and keep each other accountable.

TODD JOHNSON

It's important to have diverse ideas and viewpoints in the room. Diversity is partly how we look, but it's also how we think.

James Meadow, AI and Machine Learning Principal at Slalom

Q: How do you balance privacy with innovation and with business objectives? How do you avoid bias in AI?

Meadow: In terms of bias, it's important to have diverse ideas and viewpoints in the room. Diversity is partly how we look, but it's also how we think. Every time a story about an AI going out of control hits the front page, it seems obvious in hindsight what was going to happen. In the moment, however, it didn't seem like the obvious outcome to the team working on it. As strange as it sounds, teams really need to put themselves in the shoes of someone with the worst possible intentions – even if just for a few minutes – to make sure that they can see all the possibilities.

Unfortunately, the tech industry is not well known for diversity. In fact, it's known for quite the opposite. We all have to start doing a better job at making sure that not only do our teams look diverse, but that they think in diverse ways.

Mescall: One of the things we're finding is that on the surface, companies are taking data privacy very seriously. They'll set out policies and procedures, do everything by the book as it's drawn up, then declare "the data is now private." Then what happens is that the company – maybe because they are trying to use data to make some decisions, maybe they want to use the data to do something other than what it was intended for – takes that "private data" and moves it to a differ-

ent database. And while data may have been encrypted previously, customer or patient data is now sitting in a database that is not encrypted at all. This phenomenon of data sprawl happens at companies mostly because companies are trying to do things with the data that we all want to do: to make business decisions. Thankfully there are new techniques – AI-based techniques, in fact – that can provide you with a landscape of all your data as a company, at which point you can look at it and ensure you're protecting the data as best you can.

Fjeld: A very smart person told me one time, "There are two ways you can protect the money in the vault. You can either make the vault impenetrable, which is impossible, or you can make the money inside the vault worthless." The latter is the better way to go. We've seen that philosophy in action with past bills that largely immunizes companies from many of the proposed regulations, as long as the data is anonymous or otherwise undecipherable, unusable or unreadable Insofar as the AI plays a role in making the money in the vault worthless, I think that would be a good thing for privacy.

Q: How do you predict the AI landscape will change over the next 10 years?

Mescall: I think over the next ten years, major software vendors will start to deliver AI embedded into their products. So, if you use, say, Salesforce or Microsoft or Oracle or

any leading software company, when you buy their technologies and applications, AI will be part and parcel to their solution. The issue with that is that AI will be provider-specific. So even though you might be utilizing AI across the board, if it's all embedded, then there will still remain other problems that may not be solved by a software-embedded product. So like any other software category there will be best of breed players, companies will have AI sprawl across their business. It's going to become more complex, rather than less complex.

Meadow: With software-embedded AI, it's going to become harder and harder to determine where the decision is being made and where the AI lives in that software. The AI is going to be abstracted toward the back of the system, so to speak, and users interacting with the program won't know exactly where the decisions are being made.

Fjeld: The regulatory issues surrounding AI aren't going away, and they'll only increasingly be on Members' radar screens. The fact that data collection and use of that data is only going to accelerate over the next 10 years means that AI is going to remain a hot topic for Congress.