What happens when you feed biased training data to a machine is predictable: biased results. Bias in AI systems often stems from inherent human bias. When technologists build systems around their own experience -- even when Silicon Valley has a notable diversity problem -- or when they use training data that has had human bias involved historically, the data tends to reflect the lack of diversity or systemic bias.

Identify your projects based on the AI value chain

Some of these AI technologies can have ethical implications.

Because of this, systems inherit this bias and start to erode the trust of users. Companies are starting to realize that if they plan to gain adoption of their AI systems and realize ROI, those AI systems must be trustworthy. Without trust, they won't be used, and then the AI investment will be a waste.

Companies are combating inherent data bias by implementing programs to not only broaden the diversity of their data sets, but also the diversity of their teams. More diversity on teams enables a diverse group of people to feed systems different data points from which to learn. Organizations like AI4ALL are helping enterprises meet both of these anti-bias goals.

More human-like bots raise stakes for ethical AI use

At Google's I/O event earlier this month, the company demoed Google Duplex, an experimental Google voice assistant that was shown via a prerecorded interaction of the system placing a phone call to a hair salon on a human agent's behalf. The system did a reasonable enough job impersonating a human, even adding umms and mm-hmms, that the human on the other side was suitably fooled into thinking she was talking to another human.

This demo raised a number of significant and legitimate ethical issues of AI use. Why did the Duplex system try to fake being human? Why didn't it just identity itself as a bot upfront? Is it OK to fool humans into thinking they're talking to other humans?

Putting bots like this out into the real world where they pretend to be human, or even pretend to take over the identity of an actual human, can be a big problem. Humans don't like being fooled. There's already significant erosion in trust in online systems with people starting to not believe what they read, see or hear.

With bots like Duplex on the loose, people will soon stop believing anyone or anything they interact with via phone. People want to know who they are talking to. They seem to be fine with talking to humans or bots as long as the other party truthfully identifies itself.

Ethical AI is needed for broad AI adoption

Many in the industry are pursuing the creation of a code of ethics for bots to address potential issues, malicious or benign, that could arise, and to help us address them now before it's too late. This code of ethics wouldn't just address legitimate uses of bot technology, but also intentionally malicious uses of voice bots.

instructing the tool to ask a parent to pick up their sick child at school in order to get them out of their house so a criminal can come in while they aren't home and rob them. Bot calls from competing restaurants could make fake reservations, preventing actual customers from getting tables.

Also concerning are information disclosure issues and laws that are not up to date to deal with voice bots. For example, does it violate HIPAA laws for bots to call your doctor's office to make an appointment and ask for medical information over the phone?

Forward-thinking companies see the need to create AI systems that address ethics and bias issues, and are taking active measures now. These enterprises have learned from previous cybersecurity issues that addressing trust-related concerns as an afterthought comes at a significant risk. As such, they are investing time and effort to address ethics concerns now before trust in AI systems is eroded to the point of no return. Other businesses should do so, too.

Addressing the Ethical Issues of AI is Key to Effective Use

Enterprises must confront the ethical implications of AI use as they increasingly roll out technology that has the potential to reshape how humans interact with machines

Many enterprises are exploring how AI can help move their business forward, save time and money, and provide more value to all their stakeholders. However, most companies are missing the conversation about the ethical issues of AI use and adoption.

Even at this early stage of AI adoption, it's important for enterprises to take ethical and responsible approaches when creating AI systems because the industry is already starting to see backlash against AI implementations that play loose with ethical concerns.

For example, Google recently saw pushback with its Google Duplex release that seems to show AI-enabled systems pretending to be humans. Microsoft saw significant issues with its Tay bot that started going off the rails. And, of course, who can ignore what Elon Musk and others are saying about the use of AI.

Yet enterprises are already starting to pay attention to the ethical issues of AI use. Microsoft, for example, has created the AI and Ethics in Engineering and Research Committee to make sure the company's core values are included in the AI systems it creates.

How AI systems can be biased

AI systems can quickly find themselves in ethical trouble when left inadequately supervised. One notable example was Google's image recognition tool mistakenly classifying black people as gorillas, and the aforementioned Tay chatbot becoming a racist, sexist bigot.

How could this happen? Plainly put, AI systems are only as good as their training data, and that training data has bias. Just like humans, AI systems need to be fed data and told what that data is in order to learn from it.