Deconstructing AI

Despite the increasingly pervasive nature of AI in today's world, ambiguity abounds around basic concepts.

At its core, AI is machine learning. All applications of AI that you hear about today use machine learning (ML) in some form. Whether you talk about neural networks or deep learning, you can confidently label it ML – essentially algorithms and statistical models that are used in computer systems to perform operations without explicit programming instructions.

AI is the super-set, with most of the space within it being taken up by ML. Any area that remains is a matter of philosophical debate – interesting, but not necessarily pragmatic. Rules engines or expert systems are considered by some as AI, while others argue that they are not intelligent, as the outputs are clearly deterministic. It is simpler to think of all the AI systems in use today as the user-friendly manifestation of arcane ML models.

Learning requires self-improvement

There is considerable literature on what AI really means. It will make life easier for everybody if we accept the base condition for any AI system to be the ability to learn without explicit external intervention. From Skynet to Legion to Rekognition, the intelligence comes from being able to get better through "experience."

As long as there is an element of self-improvement baked into a system or application, it is appropriate to use the AI rubric. The learning itself can be supervised, unsupervised, semi-supervised (usually used for discerning patterns that are difficult to see), or reinforcement. But the model needs to explicitly allow for self-learning and improvement.

Let's take forecasting as an example. You can generate and execute a forecasting model in Excel at a point of time with a relatively small data set and a few variables using regression. That does not seem much like AI to an outsider, even though regression is technically an ML algorithm.

So why not set the bar a sliver higher? You can potentially enrich the ML-based model to take even more features (variables) into account for the regression model and crunch large quantities of structured as well as unstructured data, and keep running and improving the prediction over time with every single execution as more data pours in.

These distinctions are important, more so in areas where even some seasoned practitioners erroneously assume AI is always involved. RPA and chatbots are notable instances – you can enhance both with AI, but you don't need it for delivering base functionality. That's why terms like Intelligent RPA and Conversational Bots have a place in contemporary vocabulary.

AI plays with loaded dice

Any output or prediction AI provides has a probability of occurrence attached to it, often less than 100%. This is much like how the real world works ("I am 90% sure that car was gray"), only AI makes it possible to accurately quantify the probability of its predictions. As a corollary, a change in data inputs could alter the prediction provided by an AI system.

For instance, a facial-recognition match may come with a 95% or 99% accuracy level. Either way, you can get the probability attached to the match. Best to always ask your data scientist for it. The ability to quantify the uncertainty with precision is a more significant AI function than speed or scale. And all it does is draw on straightforward statistics.

Ethics is bigger than AI

Ethics is fast taking center-stage with many initiatives like explainable AI squarely aimed at bringing ethics into the equation. You can try to manage inherent biases in ML algorithms but, much as with any law, most ethical considerations need to be embedded in the ecosystem itself.

Whether it is through agreed-to global standards (with more sophisticated versions of Asimov's three laws of robotics underpinning all development), or through the unique culture of each organization that develops an AI application, our personal and community ethics will manifest themselves in how we create and use AI. As a responsible corporate citizen, there are subtle ways of explicitly ingraining ethical standards into AI; a simple control suggested by some practitioners is to use diverse developer teams that will be more attuned to reducing bias.

The bigger issue is that ethical implications are not even considered. It is a matter of time before more people start asking how a self-driving car will choose between hitting an old lady, a young child, and a litter of adorable puppies (given that those are the only options). Will a different car make a different choice? Questions – like are we willing to enhance object recognition with facial recognition for autonomous cars? – will drag higher-order issues into the debate. While a human being would find it very difficult to choose between crashing into a Nobel laureate working on world peace versus an escaped convict, technology may indeed be able to make the choice.

But should it?

By its very nature, ethics in AI will be one of the hardest nuts to crack. Imagine a world where Asimov's first law of robotics could be programmed and followed by all AI applications. Now imagine AI trying to resolve the contradiction between the first law and taking a job away from a human.

The nature of true AI

Let's end on a lighter, but equally thought-provoking, topic. A question most AI experts get asked at some point is whether the grimmer fantasies around AI have any basis. While I would personally quantify the chances of an AI takeover in the foreseeable future at less than 0.1% (a good example of subjective probability), it is a charged subject.

Driving the example of autonomous cars a bit farther: they can already do a lot more than most specific-purpose AIs, but the gap between self-driving cars and anything akin to general AI is still considerable. One, there are certain pre-conditions regarding variables like traffic, route, and weather that are required for the cars to even perform their base function. Two, envision a situation where the optimal response happened to be a one-foot jump in the air. It's easy to see why this is unlikely to even make the menu of programmable options or decision set.

The scenario where a covey of hackers took command of multiple cars in *The Fate of the Furious* is far more realistic than a car taking meaningful control of its destiny. By a factor of at least 1,000.

The lasting impact of AI on humans will be more subliminal. It will manifest itself in how we adapt to the ubiquity of AI in everyday life. From subtle changes in behavior in response to screening of human productivity, to tweaks that make our resumes and online footprint more AI-friendly, to renewed metaphysical debate around "originality" in art, in time, human responses to AI-led stimuli will rewire our thinking with more far-reaching consequences than any feat of artificial intelligence.
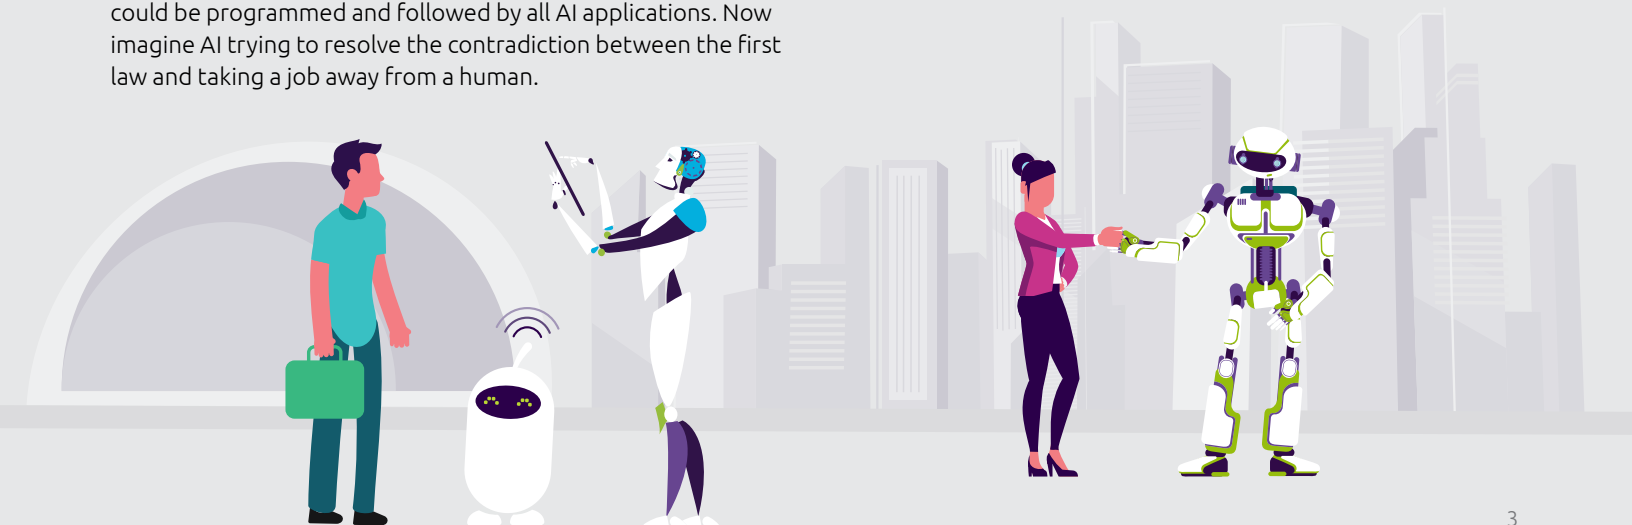
For more details contact:

Navjit Gill
Strategy and Transformation
navjit.gill@capgemini.com

About
Capgemini

A global leader in consulting, technology services and digital transformation, Capgemini is at the forefront of innovation to address the entire breadth of clients' opportunities in the evolving world of cloud, digital and platforms. Building on its strong 50-year heritage and deep industry-specific expertise, Capgemini enables organizations to realize their business ambitions through an array of services from strategy to operations. Capgemini is driven by the conviction that the business value of technology comes from and through people. It is a multicultural company of over 200,000 team members in more than 40 countries. The Group reported 2018 global revenues of EUR 13.2 billion (about $15.6 billion USD at 2018 average rate).

Learn more about us at

www.capgemini.com

People matter, results count.