



Moral AI

If we want moral AI, it's a question of nurture, not nature. **Emma Kendrew**, the Artificial Intelligence Lead for Accenture Technology in the UK and Ireland, explores how to raise systems that foster ethics, behave responsibly and are transparent to their human partners

With every advance in technology comes a wave of warnings about how it will impact society. Artificial intelligence (AI) is no different and has met a number of sceptics along the way. One of the greatest criticisms has been around ethics and how an AI can learn the right behaviours.

As humans, our own principles of right and wrong are derived from our experiences. Lessons that we are taught – either formally or informally – throughout our lives. As a result, when we talk about AI being trustworthy or ethically balanced, we should really be questioning whether it has been raised in the right way.

This is not a new dilemma: the morality of robots was in question as early on as the 1960s when Isaac Asimov came up with the 'Three Laws of Robotics'. More recently, official guidance from the British Standards Institute advises designers on how to create ethical robots. >



AI Needs to be Nurtured Like a Child

As AI becomes a more central and important factor in the lives of humans, how it's designed to make decisions becomes crucial. We believe that businesses must raise AI systems to act responsibly – and there are a number of lessons to be learned from how we educate children. Ethical constructs need to come before, not after, developing other skills. We teach children morality before maths. When they can be part of a social environment, we teach them language skills and reasoning. All of this happens before they enter a formal classroom.

Four out of five executives see AI working next to humans in their organisations as a co-worker within the next two years. It's imperative that we learn to nurture AI to address many of the same challenges faced in human education: fostering an understanding of right and wrong, and what it means to behave responsibly.

AI Needs to be Raised to Benefit Business and Society

AI is becoming smarter and more capable than ever before. With neural networks giving AI the ability to learn, the technology is evolving into an independent problem solver.

Subsequently, we need to create learning-based AI that fosters ethics and behaves responsibly – imparting knowledge without bias. Consequently, AI will be able to operate more effectively in the context of its situation. It will also be able to adapt to new requirements

“We must ensure that the AI’s human peers are trustworthy”

based on feedback from both artificial and human peers. This feedback loop is an essential, and also fundamental, part of human learning.

AI Needs to be Transparent

Employees, customers and society in general need to understand the principles used to make AI-based decisions by their organisations – especially when the stakes are high and those decisions have a direct impact on health, wealth or access to products and services.

Currently, the inner workings of AI are a mystery to the majority of people. They're even unclear to a lot of technology developers – Deepmind recently announced a goal of better understanding the decision making of its own AI.

As AI systems are primarily intended to collaborate with people, companies need to build and train them to provide clear rationale for decisions. Companies must ensure that there is transparency, so it can be clearly established whether the decision reached was fair, reliable and without bias – ultimately, was it behaving in an ethical manner?

With 72 percent of executives seeking to gain customer trust and confidence in AI by being transparent, there are

two fundamental steps we must take to achieve this. First, the data underpinning AI needs to be scrupulously clean and reliable, as any biases in data will void any subsequent intelligence and analytics.

Second, we must ensure that the AI's human peers are trustworthy and acting ethically and without bias. The ability to make socially acceptable decisions should be at the heart of every business, and if people can't trust that humans can abide by these principles, why will an AI be any different?

Fundamentally, AI that produces bad outcomes is the product of its environment and upbringing. Businesses must get ahead of the curve and go back to basics in programming AI – to make sure they're built on an unbiased, ethical grounding. If this is gotten right at the beginning, and there is a responsible and regular feedback loop, then it will become a long-term fixture. Failure to do so will see businesses unable to keep up with new regulations and the natural public demand for transparency and fairness. ■

This article was first published in Management Today.



Emma Kendrew
Artificial Intelligence Lead
Accenture Technology, UK

Previously, Emma was the Sales Lead for Accenture's Liquid Studio – a space created to facilitate innovation and speed-up software development cycles to support Accenture's clients in their digital transformation journey.

Contact Emma through:
www.criticaIeye.com