Google's release of TensorFlow could be a game-changer in the future of AI

November 13 2015, by David Tuffley

Google's machine learning software already does some pretty amazing things, such as visual translations. Credit: TensorFlow

The development of smarter and more pervasive artificial intelligence (AI) is about to shift into overdrive with the announcement by Google this week that TensorFlow, its second-generation machine-learning system, will be made available free to anyone who wants to use it.

Machine learning emulates the way the human brain learns about the

world, recognising patterns and relationships, understanding language and coping with ambiguity.

This is the technology that already provides the smarts for Google's image and speech recognition, foreign language translation and various other applications.

This is valuable technology, and it is now [open source](link); the [source code](link) is freely available and can be modified, developed in new directions and redistributed in the same way that the Linux operating system is open.

There's gold in those algorithms

There are gold-rush opportunities for imaginative commercial developers and scientists to take advantage of TensorFlow's enhanced capabilities in all kinds of ways.

For example a multi-lingual virtual assistant (VA) that anticipates your needs by using its knowledge of your patterns of daily activity combined with improved natural speech, image and pattern recognition to know what you want, when you want it, how you want it. It might also use augmented reality to overlay a real-world environment with sound, video, images or GPS information.

Beyond being an intuitive VA, TensorFlow and other such systems from Microsoft and IBM can be told to search through large data sets for something of value to you, whether it is for research purposes, business intelligence, public safety or anything else that generates data – and that's just about everything these days.

By making TensorFlow open source, Google is playing the long-game. It's positioning itself at the centre of a growing machine learning community instead of pursuing short-term profit by selling the software

or keeping it to itself. In time, any number of serendipitous developments can and will emerge from such an open community.

But Google has its work cut-out to convince the existential risk sceptics that it is still committed to its philosophy of doing business "without doing evil".

Should we be worried by advanced AI?

Intuitive applications that have an intimate place in your life will proliferate because people want them and there is much R&D effort going into getting the underlying sense-making engine to work properly. Plus, now that TensorFlow is freely available, more players will enter the game.

Some people are going to be very worried, while others will be delighted.

Microsoft founder Bill Gates and theoretical physicist Stephen Hawking have their doubts, while others such as MIT's Rodney Brooks believe that extreme AI predictions are "comparable to seeing more efficient internal combustion engines […] and jumping to the conclusion that the warp drives are just around the corner".

A brief look at history reveals a litany of doomsday warnings. We have always had the threat of asteroids, earthquakes, tsunamis, hurricanes and cyclones, plague and pestilence, drought and flooding rain. But we're still here.

Since the 1940s we had the threat of nuclear annihilation, and more recently that of catastrophic climate change. Now we can add evil (or amoral) robots to the list. Unfettered AI will become so smart that it could decide we are a plague species and should be exterminated. So say the sceptics.

Cutting through the doomsday hype, a more moderate person might recognise the need to develop AI safety protocols and risk management strategies, and get these out to industry leaders and policy makers, as suggested by the [Centre for the Study of Existential Risk](link) at Cambridge University.

The future of AI

The history of innovation clearly shows that new technology, particularly disruptive technology has a polarising effect on public opinion. It creates two camps: the optimists and the pessimists, the utopians and the distopians. Much heat is generated as the conflicting narratives do battle for dominance.

In time, the pendulum of opinion swings back and forth, the heat dissipates and extreme views moderate. The middle ground comes to be occupied by an integrated view, a compromise that has been argued out and is more or less agreed. It is a slow process though.

It will be interesting to see what impact Google's open source release of TensorFlow will have on the future of AI.

I'm going to go out on a limb and predict where the debate over strong AI will end up years from now: it is neither good nor bad but simply a tool that can be applied in countless useful ways, provided it has been developed with legally enforceable safety protocols that prevent or limit the harm that can be done through its use.

There are plenty of precedents for how we safely use potentially dangerous technology in everyday life, such as motor vehicles.

We will not wake up one morning to find our trusted household robot in the act of cutting our throat, nor will an overzealous AI factory controller

decide to eliminate all human workers for the sake of efficiency.

This story is published courtesy of [The Conversation](#) *(under Creative Commons-Attribution/No derivatives).*

Source: The Conversation