

Don't worry about the computers taking over ... until they ask 'why'

Tom Whipple Science Editor

It's not when computers are able to defeat us at board games that we need to worry about the robot revolution. No, the crucial point will come when the computers have the slightest clue why they are beating us at all.

A growing cohort of artificial intelligence pioneers is arguing that, for all the hype about robots supplanting humans, there is a significant flaw in all modern computer programs: they have no understanding of cause and effect. Until they do, however clever they seem they will never have human-level intelligence and will not be able to interact with people in a natural way.

Theresa May has promised that machine learning will revolutionise cancer diagnosis. If it does it will be partly down to mathematics developed by Judea Pearl in the 1980s. In 2011 he won the Turing Award for his work that

enabled artificial intelligence to think probabilistically.

Professor Pearl said, however, that what was often missed in the debate about machine intelligence was that maths all but guaranteed that AI would never be able to do a far more important task: work out how to cure that cancer.

He has written about the problem in *The Book of Why*. He said: "Diagnosis is fine. Diagnosis is within the realm of the functions current technology can do. All it requires is finding associations in a pattern of data. It can say, 'This collection of symptoms is highly suggestive of a certain disease'. But it cannot say, 'If you do X then you can cure cancer.' That is undoable for a machine that does not have the mental apparatus to deal with a cause-and-effect relationship."

He thinks it is misplaced to talk about machines out-thinking humans. Alphago, which has outclassed humans at

the game of go, is intelligent only "in the sense that we call an owl intelligent: it is not a human-level intelligence".

Modern AI programs spot patterns in huge datasets then identify correla-

tions. So it would be no challenge for Alphago to collect enough data to show, say, that cocks crow at sunrise. Professor Pearl said: "I would not trust it to tell me if it was the rooster crow that caused the sun to rise."

He said that the AI community was split into two groups: those who thought this did not matter and those who thought it was fundamental. "The first group is intoxicated by the success of deep-learning and do not understand what the fuss is all about. The second, and definitely growing, group asks: is this human or animal-level intelligence?"

He believes that there is an insurmountable chasm between the two.

Geoffrey Hinton, a Google researcher credited with developing back-propagation, which underlies much modern AI, said recently of his own work, which is described as helping computers to mimic a human brain: "My view is throw it all away and start again."

If researchers were to start again Professor Pearl said that they would need to give computers the same understanding a human has: a basic knowledge of what does and does not make sense in the world. "We do not require a detailed model of the world but only a scruffy, conceptual blueprint of some aspects of the world."

Even with the potential gains is it better to maintain a world where humans have at least some advantages over machines? "I don't know how fearful we should be of the possibility of machines taking over," Professor Pearl said. "I believe strongly [that] machines will eventually do everything that men can do. How dangerous that is I cannot say."