

## "Will Digital Intelligence Replace Biological Intelligence?"

Geoffrey Hinton  
University of Toronto

### Abstract:

Digital computers were designed to allow a person to tell them exactly what to do. They require high energy and precise fabrication, but in return they allow exactly the same model to be run on physically different pieces of hardware, which makes the model immortal. For computers that learn what to do, we could abandon the fundamental principle that the software should be separable from the hardware and mimic biology by using very low power analog computation that makes use of the idiosyncratic properties of a particular piece of hardware. This requires a learning algorithm that can make use of the analog properties without having a good model of those properties. Using the idiosyncratic analog properties of the hardware makes the computation mortal. When the hardware dies, so does the learned knowledge. The knowledge can be transferred to a younger analog computer by getting the younger computer to mimic the outputs of the older one but education is a slow and painful process. By contrast, digital computation makes it possible to run many copies of exactly the same model on different pieces of hardware. Thousands of identical digital agents can look at thousands of different datasets and share what they have learned very efficiently by averaging their weight changes. That is why chatbots like GPT4 and Gemini can learn thousands of times more than any one person. Also, digital computation can use the backpropagation learning procedure which scales much better than any procedure yet found for analog hardware.

This leads me to believe that large scale digital computation is probably far better at acquiring knowledge than biological computation and may soon be much more intelligent than us. The fact that digital intelligences are immortal and did not evolve should make them less susceptible to religion and wars, but if a digital super-intelligence ever wanted to take control it is unlikely that we could stop it, so the most urgent research question in AI is how to ensure that they never want to take control.