

<http://www.technologyreview.com/news/531146/what-it-will-take-for-computers-to-be-conscious/>

link no longer working;

Author Christopher Koch

What It Will Take for Computers to Be Conscious

Commentary in blue

In a refreshing change from the strong AI perspective on consciousness, Koch proposes that consciousness is an intrinsic property of all matter, just like mass or energy. He advocates for *Integrated Information Theory*, which holds that consciousness is an emergent property of huge numbers of interconnections, and that these connections could be mechanical as well as biological. He states: “If you were to build a computer that has the same circuitry as the brain, this computer would also have consciousness associated with it.” One cannot argue with such a truism, but let’s see anyone build such a computer, or even one element of such circuitry. *Integrated Information Theory* seems to me to attempt to explain away what David Chalmers has called “the hard problem of consciousness” by throwing up the hands and saying, “if there are enough connections, consciousness will happen”.

The world’s best-known consciousness researcher, Christof Koch, says machines could one day become self-aware. October 2, 2014

Christof Koch, chief scientific officer of the Allen Institute for Brain Science in Seattle, believes consciousness is an intrinsic property of matter, just like mass or energy. Organize matter in just the right way, as in the mammalian brain, and voilà, you can feel.

Koch, now 57, has spent nearly a quarter of a century trying to explain why, say, the sun feels warm on your face. But after writing three books on consciousness, Koch says researchers are still far from knowing why it occurs, or even agreeing on what it is. It’s a difficult problem (see “Cracking the Brain’s Codes”). That is one reason that Koch left his position at Caltech in 2011 to become part of a \$500 million project launched by the billionaire Paul Allen, Microsoft’s cofounder.

The Allen Institute’s goal is to build a detailed atlas of every neuron and synapse in the mammalian brain. That would give neuroscience a firehose of data similar to what the Human Genome Project achieved.

But Koch hasn’t given up his search for a grand theory that could explain it all. In fact, he thinks consciousness could be explained by something called “integrated information theory,” which asserts that consciousness is a product of structures, like the brain, that can both store a large amount of information and have a critical density of interconnections between their parts.

To Koch, the theory provides a means to assess degrees of consciousness in people with brain damage, in species across the animal kingdom, and even, he says, among machines.

So, you think a computer can be conscious?

I gave a lecture [last week] at MIT about Integrated Information Theory, developed by Giulio Tononi at the University of Wisconsin. This is a theory that makes a very clear prediction: it says that consciousness is a property of complex systems that have a particular “cause-effect” repertoire. They have a particular way of interacting with the world, such as the brain does, or in principle, such as a computer could. If you were to build a computer that has the same circuitry as the brain, this computer would also have consciousness associated with it. It would feel like something to be this computer. However, the same is not true for digital simulations.

If I build a perfect software model of the brain, it would never be conscious, but a specially designed machine that mimics the brain could be?

Correct. This theory clearly says that a digital simulation would not be conscious, which is strikingly different from the dominant functionalist belief of 99 percent of people at MIT or philosophers like Daniel Dennett. They all say, once you simulate everything, nothing else is required, and it’s going to be conscious.

I think consciousness, like mass, is a fundamental property of the universe. Simulating something is not the real thing.

It’s the same thing with consciousness. In 100 years, you might be able to simulate consciousness on a computer. But it won’t experience anything.

I am not saying consciousness is a magic soul. It is something physical. Consciousness is always supervening onto the physical. But it takes a particular type of hardware to instantiate it. A computer made up of transistors, moving charge on and off a gate, with each gate being connected to a small number of other gates, is just a very different cause-and-effect structure than what we have in the brain, where you have one neuron connected to 10,000 input neurons and projecting to 10,000 other neurons. But if you were to build the computer in the appropriate way, like a neuromorphic computer [see “Thinking in Silicon”], it could be conscious.