

A Theory of Reality as More Than the Sum of Its Parts

Using the mathematical language of information theory, Hoel and his collaborators claim to show that new causes — things that produce effects — can emerge at macroscopic scales. They say coarse-grained macroscopic states of a physical system (such as the psychological state of a brain) can have more causal power over the system's future than a more detailed, fine-grained description of the system possibly could. Macroscopic states, such as desires or beliefs, "are not just shorthand for the real causes," explained [Simon DeDeo](#), an information theorist and cognitive scientist at Carnegie Mellon University and the Santa Fe Institute who is not involved in the work, "but it's actually a description of the real causes, and a more fine-grained description would actually miss those causes."

"To me, that seems like the right way to talk about it," DeDeo said, "because we do want to attribute causal properties to higher-order events [and] things like mental states."

Hoel and collaborators have been developing the mathematics behind their idea since 2013. In [a May paper](#) in the journal *Entropy*, Hoel placed causal emergence on a firmer theoretical footing by showing that macro scales gain causal power in exactly the same way, mathematically, that [error-correcting codes](#) increase the amount of information that can be sent over information channels. Just as codes reduce noise (and thus uncertainty) in transmitted data — Claude Shannon's 1948 insight that formed the bedrock of information theory — Hoel claims that macro states also reduce noise and uncertainty in a system's causal structure, strengthening causal relationships and making the system's behavior more deterministic.

"I think it's very significant," [George Ellis](#), a South African cosmologist who has also [written about top-down causation](#) in nature, said of Hoel's new paper. Ellis thinks causal emergence could account for many emergent phenomena such as [superconductivity](#) and topological phases of matter. Collective systems like bird flocks and superorganisms — and even simple structures like crystals and waves — might also exhibit causal emergence, researchers said.

The work on causal emergence is not yet widely known among physicists, who for centuries have taken a reductionist view of nature and largely avoided further philosophical thinking on the matter. But at the interfaces between physics, biology, information theory and philosophy, where puzzles crop up, the new ideas have generated excitement. Their ultimate usefulness in explaining the world and its mysteries — including consciousness, other kinds of emergence, and the relationships between the micro and macro levels of reality — will come down to whether Hoel has nailed the notoriously tricky notion of causation: Namely, what's a cause? "If you

brought 20 practicing scientists into a room and asked what causation was, they would all disagree,” DeDeo said. “We get mixed up about it.”

A Theory of Cause

In a fatal drunk driving accident, what’s the cause of death? Doctors name a ruptured organ, while a psychologist blames impaired decision-making abilities and a sociologist points to permissive attitudes toward alcohol. Biologists, chemists and physicists, in turn, see ever more elemental causes. “Famously, Aristotle had a half-dozen notions of causes,” DeDeo said. “We as scientists have rejected all of them except things being in literal contact, touching and pushing.”

The true causes, to a physicist, are the fundamental forces acting between particles; all effects ripple out from there. Indeed, these forces, when they can be isolated, appear perfectly deterministic and reliable — physicists can predict with high precision the outcomes of particle collisions at the Large Hadron Collider, for instance. In this view, causes and effects become hard to predict from first principles only when there are too many variables to track.

Furthermore, philosophers have argued that causal power existing at two scales at once would be twice what the world needs; to avoid double-counting, the “exclusion argument” says all causal power must originate at the micro level. But it’s almost always easier to discuss causes and effects in terms of macroscopic entities. When we look for the cause of a fatal car crash, or Romeo’s decision to start climbing, “it doesn’t seem right to go all the way down to microscopic scales of neurons firing,” DeDeo said. “That’s where Erik [Hoel] is jumping in. It’s a bit of a bold thing to do to talk about the mathematics of causation.”

Friendly and large-limbed, Hoel grew up reading books at Jabberwocky, his family’s bookstore in Newburyport, Massachusetts. He studied creative writing as an undergraduate and planned to become a writer. (He still [writes fiction](#) and has started a novel.) But he was also drawn to the question of consciousness — what it is, and why and how we have it — because he saw it as an immature scientific subject that allowed for creativity. For graduate school, he went to Madison, Wisconsin, to work with Tononi — the only person at the time, in Hoel’s view, who had a truly scientific theory of consciousness.

Tononi conceives of consciousness as information: bits that are encoded not in the states of individual neurons, but in the complex networking of neurons, which link together in the brain into larger and larger ensembles. Tononi argues that this special “integrated information” corresponds to the unified, integrated state that we experience as subjective awareness. Integrated information theory has gained prominence in the last few years, even as debates have ensued about whether it is an

accurate and sufficient proxy for consciousness. But when Hoel first got to Madison in 2010, only the two of them were working on it there.

Tononi tasked Hoel with exploring the general mathematical relationship between scales and information. The scientists later focused on how the amount of integrated information in a neural network changes as you move up the hierarchy of spatiotemporal scales, looking at links between larger and larger groups of neurons. They hoped to figure out which ensemble size might be associated with maximum integrated information — and thus, possibly, with conscious thoughts and decisions. Hoel taught himself information theory and plunged into the philosophical debates around consciousness, reductionism and causation.

Hoel soon saw that understanding how consciousness emerges at macro scales would require a way of quantifying the causal power of brain states. He realized, he said, that “the best measure of causation is in bits.” He also read the works of the computer scientist and philosopher [Judea Pearl](#), who developed a logical language for studying causal relationships in the 1990s called causal calculus. With Albantakis and Tononi, Hoel formalized a measure of causal power called “effective information,” which indicates how effectively a particular state influences the future state of a system. (Effective information can be used to help calculate integrated information, but it is simpler and more general and, as a measure of causal power, does not rely on Tononi’s other ideas about consciousness.)

The researchers showed that in simple models of neural networks, the amount of effective information increases as you coarse-grain over the neurons in the network — that is, treat groups of them as single units. The possible states of these interlinked units form a causal structure, where transitions between states can be mathematically modeled using so-called Markov chains. At a certain macroscopic scale, effective information peaks: This is the scale at which states of the system have the most causal power, predicting future states in the most reliable, effective manner. Coarse-grain further, and you start to lose important details about the system’s causal structure. Tononi and colleagues hypothesize that the scale of peak causation should correspond, in the brain, to the scale of conscious decisions; based on brain imaging studies, Albantakis guesses that this might happen at the scale of neuronal microcolumns, which consist of around 100 neurons.

Causal emergence is possible, Hoel explained, because of the randomness and redundancy that plagues the base scale of neurons. As a simple example, he said to imagine a network consisting of two groups of 10 neurons each. Each neuron in group A is linked to several neurons in group B, and when a neuron in group A fires, it usually causes one of the B neurons to fire as well. Exactly which linked neuron fires is unpredictable. If, say, the state of group A is {1,0,0,1,1,1,0,1,1,0}, where 1s and 0s represent neurons that do and don’t fire, respectively, the resulting state of group B

can have myriad possible combinations of 1s and 0s. On average, six neurons in group B will fire, but which six is nearly random; the micro state is hopelessly indeterministic. Now, imagine that we coarse-grain over the system, so that this time, we group all the A neurons together and simply count the total number that fire. The state of group A is {6}. This state is highly likely to lead to the state of group B also being {6}. The macro state is more reliable and effective; calculations show it has more effective information.

A real-world example cements the point. “Our life is very noisy,” Hoel said. “If you just give me your atomic state, it may be totally impossible to guess where your future [atomic] state will be in 12 hours. Try running that forward; there’s going to be so much noise, you’d have no idea. Now give a psychological description, or a physiological one: Where are you going to be in 12 hours?” he said (it was mid-day). “You’re going to be asleep — easy. So these higher-level relationships are the things that seem reliable. That would be a super simple example of causal emergence.”

For any given system, effective information peaks at the scale with the largest and most reliable causal structure. In addition to conscious agents, Hoel says this might pick out the natural scales of rocks, tsunamis, planets and all other objects that we normally notice in the world. “And the reason why we’re tuned into them evolutionarily [might be] because they are reliable and effective, but that also means they are causally emergent,” Hoel said.

Brain-imaging experiments are being planned in Madison and New York, where Hoel has joined the lab of the Columbia neuroscientist [Rafael Yuste](#). Both groups will examine the brains of model organisms to try to home in on the spatiotemporal scales that have the most causal control over the future. Brain activity at these scales should most reliably predict future activity. As Hoel put it, “Where does the causal structure of the brain pop out?” If the data support their hypothesis, they’ll see the results as evidence of a more general fact of nature. “Agency or consciousness is where this idea becomes most obvious,” said William Marshall, a postdoctoral researcher in the Wisconsin group. “But if we do find that causal emergence is happening, the reductionist assumption would have to be re-evaluated, and that would have to be applied broadly.”

New Philosophical Thinking

[Sara Walker](#), a physicist and astrobiologist at Arizona State University who studies the origins of life, hopes measures like effective information and integrated information will help define what she sees as the gray scale leading between nonlife and life (with viruses and cell cycles somewhere in the gray area). Walker has been collaborating with Tononi’s team on studies of real and artificial cell cycles, with preliminary indications that integrated information might correlate with being alive.

In [other recent work](#), the Madison group has developed a way of measuring causal emergence called “black-boxing” that they say works well for something like a single neuron. A neuron isn’t simply the average of its component atoms and so isn’t amenable to coarse-graining. Black-boxing is like putting a box around a neuron and measuring the box’s overall inputs and outputs, instead of assuming anything about its inner workings. “Black-boxing is the truly general form of causal emergence and is especially important for biological and engineering systems,” Tononi said in an email.

Walker is also a fan of Hoel’s new work tracing effective information and causal emergence to the foundations of information theory and Shannon’s noisy-channel theorem. “We’re in such deep conceptual territory it’s not really clear which direction to go,” she said, “so I think any bifurcations in this general area are good and constructive.”

[Robert Bishop](#), a philosopher and physicist at Wheaton College, said, “My take on EI” —effective information — “is that it can be a useful measure of emergence but likely isn’t the only one.” Hoel’s measure has the charm of being simple, reflecting only reliability and the number of causal relationships, but according to Bishop, it could be one of several proxies for causation that apply in different situations.

Hoel’s ideas do not impress [Scott Aaronson](#), a theoretical computer scientist at the University of Texas, Austin. He says causal emergence isn’t radical in its basic premise. After reading Hoel’s recent essay for the Foundational Questions Institute, “Agent Above, Atom Below” (the one that featured Romeo and Juliet), Aaronson said, “It was hard for me to find anything in the essay that the world’s most orthodox reductionist would disagree with. Yes, of course you want to pass to higher abstraction layers in order to make predictions, and to tell causal stories that are predictively useful — and the essay explains some of the reasons why.”

It didn’t seem so obvious to others, given how the exclusion argument has stymied efforts to get a handle on higher-level causation. Hoel says his arguments go further than Aaronson acknowledges in showing that “higher scales have provably more information and causal influence than their underlying ones. It’s the ‘provably’ part that’s hard and is directly opposite to most reductionist thinking.”

Moreover, causal emergence isn’t merely a claim about our descriptions or “causal stories” about the world, as Aaronson suggests. Hoel and his collaborators aim to show that higher-level causes — as well as agents and other macroscopic things — ontologically exist. The distinction relates to one that the philosopher David Chalmers makes about consciousness: There’s the “easy problem” of how neural circuitry gives rise to complex behaviors, and the “hard problem,” which asks, essentially, what distinguishes conscious beings from lifeless automatons. “Is EI measuring causal power of the kind that we feel that we have in action, the kind that we want our

conscious experiences or selves to have?" said Hedda Hassel Mørch, a philosopher at New York University and a protégé of Chalmers'. She says it's possible that effective information could "track real ontological emergence, but this requires some new philosophical thinking about the nature of laws, powers and how they relate."

The criticism that hits Hoel and Albantakis the hardest is one physicists sometimes make upon hearing the idea: They assert that noise, the driving force behind causal emergence, doesn't really exist; noise is just what physicists call all the stuff that their models leave out. "It's a typical physics point of view," Albantakis said, that if you knew the exact microscopic state of the entire universe, "then I can predict what happens until the end of time, and there is no reason to talk about something like cause-effect power."

One rejoinder is that perfect knowledge of the universe isn't possible, even in principle. But even if the universe could be thought of as a single unit evolving autonomously, this picture wouldn't be informative. "What is left out there is to identify entities — things that exist," Albantakis said. Causation "is really the measure or quantity that is necessary to identify where in this whole state of the universe do I have groups of elements that make up entities? ... Causation is what you need to give structure to the universe." Treating causes as real is a necessary tool for making sense of the world.

Maybe we sort of knew all along, as Aaronson contends, that higher scales wrest the controls from lower scales. But if these scientists are right, then causal emergence might be how that works, mathematically. "It's like we cracked the door open," Hoel said. "And actually proving that that door is a little bit open is very important. Because anyone can hand-wave and say, yeah, probably, maybe, and so on. But now you can say, 'Here's a system [that has these higher-level causal events]; prove me wrong on it.'"

Correction: This article was revised on June 1, 2017, to clarify that Erik Hoel was not Giulio Tononi's first collaborator on integrated information theory and that Tononi, Hoel and Larissa Albantakis formalized, but did not devise, the measure used to assess causal emergence.