Science beyond the Horizon with Stephen Wolfram (10.11.2020)

Q&A from the chat conversation

Orion Maxted: If all humans are computationally equivalent, why do some appear to be smarter than others?

Orion Maxted: Where does the threshold of computational equivalence arise in the animal kingdom? Are humans at threshold? Or are apes/mice/ants/etc as smart as us? Where do you see the threshold?

Jonathan Gorard (to Orion): one of the implications of the principle of computational equivalence is that there can't really be an abstract definition of intelligence - it's ultimately bound up with human culture and history. In other words, when one describes a person (or an animal) as being "smarter" than another person or animal, one isn't really saying that the computations they're performing are inherently more sophisticated (since most such computations are ultimately equivalent); rather one is saying that the computations have a great degree of "fit" with particularly activities (such as science, philosophy, art, etc.) that we, for purely cultural reasons, associate with "intelligence".

George Davidescu: Is looking for useful rules in the computational universe like searching for winning lottery numbers? Are there any salient features in the rules (the output) that can help direct our search?

Chris Verzijl: Seconding a previous question: Were the hypergraph experiments randomly found (lottery tickets), or was there some guiding principle in the search in order to generate interesting geometry and dimensionality?

Jonathan Gorard (to George and Chris): It's mostly a lottery. We have a few heuristics that we use for quickly determining whether a given rule is likely to be interesting or not (e.g. we have algorithms for testing whether any non-trivial interactions can occur between different hypergraph regions, or whether a given rule is just a simple reflection or translation of a rule that we've already seen, etc.), but for the most part we just have to run things and find out. There's a popular saying in the project (originally coined by Stephen): "The computational animals are always smarter than you are!" (An example of such a heuristic: if we can prove that the causal network, i.e. the pattern of causal relationships between events, is always a tree, then that's an indication that the rule is really just a neighbourhood-independent substitution system, and hence probably not particularly interesting. But, as with all heuristics, it's not perfect…)

Anonymous: I think Stephen refers to computational irreducibility (CI). If both A and B are CI then they are equivalent. If you can't tell for rule 30 where it will end up for T -> infinity and the brain would be thought to also contain CI they are equivalent.

Abdelmonim Artoli: If physics is just a use case for building our computational universe, and despites its complex development in understanding our universe, yet failing, would you highlight on how would we compute fundamentally to obtain a rule that governs a certain phenomenon? Would not physics be our universal Turing machine to which all hypergraphs and hyper-computers relax as from Church-Turing thesis?

Jonathan de Bouter: I have a question about the rule-mining process: I'm curious as to what the check/evaluation looks like that evaluates if a given rule's emergent hypergraph nicely corresponds to physics theories/rules. I'm hoping a non-exact characterization for laypeople can be given as an answer, that something can be said about it without going into the difficult exact detailed stuff too much.

Jonathan Gorard (to Jonathan): That's a great question Jonathan! We have a bunch of automated tests to determine whether a rule is likely to be compatible with special relativity, general relativity, quantum mechanics, etc. One simple example is causal invariance (which is known to be necessary condition for special relativity), which we can test for by determining whether the causal network is always the same for each updating order (the function for testing that is here https://resources.wolframcloud.com/FunctionRepository/resources/CausalInvariantQ)
Another example is testing to see whether the leading-order term in the growth rate of geodesic balls/cones in the hypergraph/causal network grows too rapidly (if it does, that is an indication that the causal network does not converge to a finite-dimensional manifold in the continuum limit, which we know is incompatible with general relativity).

Gerard 't Hooft: It is particularly difficult to reproduce Lorentz invariance in a CA. Should be possible, but how?

Orsan Senalp: How would you think of particle entanglement from this perspective?

Renske Vroomans: If we need a hypothetical "head" of the Turing machine to update the state of the universe, what makes it move? Is that outside of time? Is it within the rules?

Gerard 't Hooft: If you want to apply such ideas to the real world, there are two problems: 1) the real world looks very different; we have is a universal configuration called "vacuum" and there is a very remarkable set of "elementary particles" that move around following rules we call "quantum mechanics", QFT and so on. Can such structures be reproduced in a CA? This leads to 2) Do you have any `strategy' to approach and compare our world with a CA, so as to answer questions like "where does the electron come from? And so on.

Jonathan Gorard (to Gerard): why are you asking about CAs? We consider hypergraph models over CAs for precisely the reasons that Stephen mentioned earlier (CA topology is far too rigid, it's almost impossible to make them diffeomorphism invariant, etc.)

If you're asking about how Lorentz symmetry is recovered in the case of these hypergraph models (definitely not CAs…), then that's discussed at length in our various papers. Stephen has a rather nice demonstration in his paper, and I give a general proof (that causal invariance implies diffeomorphism invariance) here https://arxiv.org/abs/2004.14810, and in other places.

Anonymous : Lorentz transformation with speed of light replaced by some sort of processor speed?

Anonymous : What's the distance between each "space point"? Is this distance constant?

Izabelė Jonušaitė (to Anonymous): I believe the relations between space points are not spatial, space and time are claimed to emerge from them

Jonathan Gorard (to Anonymous, Izabele): Izabele is correct! The spatial metric tensor is recovered from the geodesic distance between vertices in the hypergraph (i.e. the minimum number of hyperedges that one needs to traverse to get from one vertex to the other); the vertices do not have any a priori distance defined between them.

Anonymous (to Izabele): Thanks. So the question would be, how many "space points" would exist inside an electron (or some measurement of space? Could the number of "space points" inside the space of an electron be infinite?

Jonathan Gorard (to Anonymous): Our (very preliminary!) dimensional analysis suggests that a single electron may have as many as 10^60 elementary hyperedges within it: https://www.wolframphysics.org/technical-introduction/potential-relation-to-physics/units-and-scales/. But this estimate is very rough, and based on several assumptions that may be incorrect. But crucially, it's not infinite!

Emiel Robben: Have your models led to a sort of precise results yet? Einstein's equations for example lead to very precise predictions of things in nature.

Jonathan Gorard (to Emiel): Can you define what you mean by a precise result? If you mean "well-defined mathematical predictions about nature", then yes, we have several! Just a few weeks ago we made a collection of new predictions about quantum information theory (see, e.g. https://arxiv.org/abs/2010.02752).

Eric Lorenz: while the dimensions of a CA or nodes in a hypergraph do not necessarily need to directly correspond to space or time, they all assume some sort of computational substrate or (discretized) ether. Galilean invariance remains a big challenge in such systems since Konrad Zuse's Calculating Space (and earlier). What are your thoughts here? Is Galilean invariance a phenomenon emerging from simpler rules?

Jonathan Gorard [21:09:06]: The paper that SW is referring to https://arxiv.org/abs/2010.02752

Chris Verzijl: Are there any refs to papers which address the mapping to QM and GRT as their focus, and limitations Stephen is claiming as "mistakes" in 20th century physics? Also, how do those computations of GR problems like the black-hole mergers compare (tractability, runtime) vs standard techniques in computational GR actually, if I may ask?

Jonathan Gorard (to Chris): They're tend to be less efficient than standard numerical GR techniques, since normal numerical GR assumes a fairly fixed coordinate structure (e.g. curvilinear coordinates), whereas in our hypergraph model the coordinate structure is intrinsically dynamic, and that's (both algorithmically and computationally) harder to deal with.

Emiel Robben: how do your models cope with a big bang? How do your models cope with the (speculative) creation of constants in nature and laws of nature during the big bang?

Jonathan Gorard (to Emiel): What do you mean by "cope with" a Big Bang? There exist classes of models that we can prove are compatible with the conformal structure of lambda-CDM cosmology, if that's what you mean...

Emiel Robben (to Jonathan Gorard) I mean: "(how) is there given a place to the big bang in your models?"

Jonathan Gorard (to Emiel): Well, as discussed here https://arxiv.org/abs/2004.14810 and in various other places, a generic feature of these models is for them to start in effectively infinite numbers of dimensions (with arbitrarily densely-connected hypergraphs) and gradually converge down to something finite dimensional. And, as I alluded to above, that reproduces the same predictions regarding the conformal structure of spacetime as a standard inflationary modification to lambda-CDM.

Gordon Klaus: It's fascinating that general relativity and quantum mechanics seem to be two instantiations of the same abstract theory. Is this a new result? Do other attempts at unification (string theory, loop quantum gravity) give any similar indications?

Jonathan Gorard (to Gordon): The closest that we're aware of is the general principle of holography (in which lower-dimensional conformal field theories are dual to higher-dimensional gravitational theories); the correspondence that we seem to be uncovering appears to be some effective generalisation of holography.

George Davidescu: Which one is the rule for our universe? The example in the top left?

Renske Vroomans: Would the "rulial reference frame" still meaningfully apply to beings that exist at a vastly different scale than that of the universes' rule? Wouldn't "causal invariance" somehow limit the possibilities of how you can see the universe at this "higher" level?

Jonathan Gorard (to Renske): Very interesting question! We don't really know the answer - the correspondence between what physics would look like to "alien intelligence" which operate at vastly different length scales to us, and the resultant foliations that they would observe of the rulial multiway graph, is still largely mysterious...

Joanna Luc: I don't understand why any cases of cellular automata can be regarded as an example of computational irreducibility. There is a simple rule describing them, namely, the one by which we have generated them. This is different from the case of truly random spread of events, which might be truly irreducible in the sense that we cannot find any rule generating it.

Jonathan Gorard (to Joanna): But the key question is how do you distinguish between "true randomness" (whatever that means) and "artificial randomness" (e.g. as generated by a CA). Suppose I show you the output of a computer program and ask you whether it's "truly random" or simply the output of an elementary CA rule or similar. How would you tell the difference?

The point is that computational irreducibility prevents you from being able to make that determination in general with any finite amount of computational effort.

Joanna Luc: But then what IS an example of computational reducibility? The example given was that of the laws of motion in classical physics. But I don't see the difference between these laws and CA laws. They both generate the output from simple rules. But also it both cases it might take a lot of effort to derive what the actual behaviour will be given the basic rules.

Jonathan Gorard (to Joanna): That example was specifically the laws of motion for the case of a classical 2-body gravitational system. In that case, both the laws and behaviour are correspondingly simple (it is computationally easy to derive one from the other, etc.), making it a reasonable candidate for reducibility. Once you have a 3-body system, for instance, there no longer exists any such nice correspondence, making it reasonable to say that it's irreducible.

Vincenzo De Florio: a general systems theory of everything

Orion Maxted: Is there a meta-artistic space? What are its limits?

Jonathan Gorard (to Orion): It's not clear that art is finitely axiomatisable ;) (Making it hard to enumerate multiway systems…)

Amos: I was wondering are you expecting the rule underlying the universe to be unique and stable, or might the rule also evolve / diverge in different regions / be unstable globally? How could such changes be ruled out if you expect only one stable rule?

Chris Verzijl: I have a similar question to Amos, concerning rules being mapped to laws, whether they should have long horizon stability that's provable in some sense (re the law, not ruling out deterministic chaos under the law, etc), versus even the law being mutable (which I guess would be unknowable due to irreducibility)? (law = physical law in that last question)

Jonathan Gorard (to Amos and Chris): Presumably even a rule that is "unstable" will still have a higher-level rule describing how it changes with time? In which case you're just kicking the problem up one level of abstraction (and the rulial multiway system accommodates all such levels anyway).

Chris Verzijl: OK fair to some extent, but perhaps I missed the idea that that's a necessary feature of the rulial approach...

Jonathan Gorard (to Amos and Chris): (If, on the other hand, you're asking what would happen if the rule were "unstable" at all levels of abstraction, then… well… I guess the universe doesn't follow any definite rules, and theoretical science is kind of a non-starter ;))

Anonymous: So can you predict things like the position of a particle in the future or what happens at the end of the universe?

Oscar de Wit: How does the continuum limit work for causal graphs?

Alicia Castro: What seems to me is that the space(time) structure you generate is purely classical, you get Einstein's Equations, but what does this approach tell us about spacetime at very high energy (maybe, namely, quantum spacetime/gravity)

Caroline Bauer: (How) could such a theory, if it continues to succeed reproducing known physics, be used to make predictions about new theoretical concepts?

Aleksandar Shulevski: The formalism describes things we know, so in that sense it is a language. Can we use it to say for example what dark matter is, and predict what should we look for, to measure its localization on the sky?

George Davidescu: computational irreducibility reminds me of calculating irrational numbers

Emiel Robben: have your models found a new field of study in Physics yet? Quantum mechanics and relativity were discovered already before you simulated them with your model. Can you "predict" laws of nature where they have not been discovered yet?

Jonathan Gorard [22:20:04] (More details about what Stephen is currently talking about in regards to the relationship between homotopy type theory, univalence, multiway systems, etc., can be found here https://www.wolframphysics.org/bulletins/2020/08/a-candidate-geometrical-formalism-for-the-foundations-of-mathematics-and-physics/)