

Is Intelligence Non-Computational Dynamical Coupling?

JONATHAN A. SIMON
University of Montreal

Abstract: Is the brain *really* a computer? In particular, is our intelligence a computational achievement: is it because our brains are computers that we get on in the world as well as we do? In this paper I will evaluate an ambitious new argument to the contrary, developed in Landgrebe and Smith (2021a, 2022). Landgrebe and Smith begin with the fact that many dynamical systems in the world are difficult or impossible to model accurately (*inter alia*, because it is intractable to find exact solutions to the differential equations that describe them—meaning we have to approximate—but at the same time they are such that small differences in starting conditions lead to big differences in final conditions, thwarting accurate approximation). Yet we manage to survive and thrive in a world full of such systems. Landgrebe and Smith argue from these premises that it is not because our brains are computers that we get on as well as we do: instead it is because of the various ways that we dynamically couple with such systems, these couplings themselves impossible to model well enough to emulate *in silico*. Landgrebe and Smith thus defend a dynamical systems model in the lineage of Gibson (1979), Van Gelder (1995), and Thompson (2007), though their focus is on decisively refuting the computationalist alternatives rather than developing the positive account. Here I will defend the claim that human intelligence is genuinely computational (and that whole brain emulation and others forms of AGI may be possible) against this argument.

1. INTRODUCTION

Today the computer is among our most transformational technologies, and it is hard to resist the idea that the brain is one (i.e., that human-level intelligence is explainable in terms of computations that our brains implement). Not too long ago, steam engines were the most transformational technology, and many thought the brain was a steam engine. Have we arrived at the truth, or will computationalism seem as quaint in ages to come as hydraulicism does today (or, indeed, quainter)? In their recent book, Landgrebe and Smith (2022) make an ambitious case that computationalism is just a fad. On their view, while the human brain may do some computational modelling, this is not the cornerstone of our intelligence because coping in the world is computationally intractable (or perhaps even non-computable).

As I construe it, they develop two complementary arguments that our brains are not computers (i.e., that human-level intelligence is not explainable in terms of computations

that our brains implement). I'll get right to them. One bit of terminology: a model is 'adequate and synoptic' if its predictions are accurate enough to be reliable in the context for which it is needed (full definition in §3.2).¹ I'll clarify the other key terms below. Now the arguments (exact formulation mine, not theirs). The first of the two was originally presented in an earlier article, Landgrebe and Smith (2021a):

The Coping Argument

1. There is a non-empty class C of complex systems that cannot be adequately and synoptically computationally modelled.
2. We cope.
3. If our brains are computers, we cope iff we adequately and synoptically computationally model some complex systems from C (those in our environmental niche).
4. Therefore, our brains are not computers.

Landgrebe and Smith (2022) offers further support for each of the premises of the above argument, and supplements it with another argument, focused on the possibility of whole brain emulation:

The Emulation Argument

1. There is a non-empty class C of complex systems that cannot be adequately and synoptically computationally modelled.
5. Our brains are systems in class C.
6. If our brains cannot be adequately and synoptically modelled, they cannot be emulated *in silico* with the requisite fidelity for whole brain emulation.
7. If our brains cannot be emulated *in silico* with the requisite fidelity for whole brain emulation, then our brains are not computers.
8. Therefore, our brains are not computers.

These two arguments amount to a distinctive, ambitious challenge to computationalism about intelligence. How to evaluate them? There are many important questions about premise (1), but I'll mostly leave those for another day, and grant premise (1) for the sake of argument. Here, I'll challenge premises (2), (3), (6) and (7).

Here is how I'll proceed. In §2 I situate the Landgrebe and Smith argument in the literature. In §3 I will develop the key terms in the argument, and motivate premise one. In §4 I challenge the coping argument. I contend that there is a critical equivocation in the notion of 'coping'. In §4.1 I argue that if 'coping' means, *performing as well as you would if you had adequate and synoptic models of complex environmental systems*, then premise two is false: we can't cope. I'll call this the *Argument from Human Frailty*. In §4.2 I argue that if 'coping' means anything less demanding than that, then premise three is false. I'll call this the *Argument from the General Usefulness of Computation*. In §5 I challenge the emulation argument. I contend that there is a critical equivocation in the notion of 'emulated *in silico* with the requisite fidelity for whole brain emulation'. Construed as a claim on the type level, premise six is false, while construed as a claim on the token level, premise seven is false. I'll call my argument here the *Argument from Synthetic Music*. Finally in §6 I'll reflect on the morals for artificial general intelligence (AGI) and highlight the various points at which I agree with Landgrebe and Smith, and where I think further research would be helpful.

2. THE VIA NEGATIVA

The Landgrebe-Smith argument combines elements from several arguments in the literature. I'll briefly comment on some of these including arguments that cognition involves analog rather than digital computation, arguments for the dynamicist hypothesis, and the Lucas-Penrose argument.

Cognition as Analog Computation: In the relevant sense, analog computers are devices that store information in physical quantities in a continuous format and use measurements to perform computation. Analog computers (arguably) date back to antiquity, but on any construal predate digital computers. Vannevar Bush's *Differential Analyser* dates to 1931 and Shannon's *General Purpose Analog Computer* dates to 1941. One of the primary applications of these devices is the computation of differential equations. Gerard (1951) is an early defense of the claim that neural computation is analog. Maley (2018) is a more recent entry. We can construe Landgrebe and Smith as building on results from the literature on analog computation such as those in Pour-El and Richards (1974, 1981) identifying limits on the computability of solutions to some differential equations.² Crucially, Landgrebe and Smith argue that intelligence is not computation, analog or digital.

The Dynamicist Hypothesis: In the nineties, van Gelder, Port and others argued that the best model for cognition is that of a dynamical system rather than a Turing machine or a connectionist network.³ Such systems can be understood in terms of dynamical coupling: for van Gelder, the exemplar is the Watt Governor, a mechanism that attaches to a flywheel, regulating its speed by a centrifugal force, though many of the mechanisms used in analog computation also exemplify the idea. Others have defended the model by applying it, e.g. to explanations of the A-not-B error,⁴ aspects of working memory,⁵ or (in a work I co-authored) aspects of phenomenal consciousness.⁶ See Favela (2020) for a recent survey. We can construe Landgrebe and Smith (2022) as an attempt to take the further step of decisively refuting rival views.

The Lucas-Penrose Argument: According to proponents of this argument, for any sufficiently expressive, consistent formal system F, we humans have the ability to come to know the truth of the Gödel sentence G asserting that G cannot be proven in F. However, if formal systems can be said to come to know anything, it must be through proving that thing. Thus for each such system there is something that we can know but that the system itself cannot, meaning, by diagonalization, that we cannot be identical to any such system.⁷ This comes very close to showing that human intelligence involves non-computable abilities. However, it hinges on the problematic premise that for any sufficiently expressive, consistent formal system F, humans have the ability to come to know that its Gödel sentence is true. There are reasons to doubt this premise.⁸ But even if it is true, coping with complex dynamic environmental systems might still involve only computational abilities. Likewise, coping might be non-computable, though there are some truths G that we cannot know. Thus while there is a parallel between the Lucas-Penrose and the Landgrebe-Smith arguments, they are independent.

A few other arguments merit honorable mention as well: Dreyfus' (1992) argument that computers may be able to *know that* but they will never *know how* (or cope), Larson's (2021) argument that computers may be able to do induction and deduction, but they will never get abduction, or Cantwell Smith's (2019) argument that computers may be able to *reckon* but they are incapable of *judgment*. There are moments in Landgrebe and Smith (2022) that rhyme with each of these (*coping* presumably being the same notion Dreyfus has in mind, and arguably requiring both *judgment* and *abduction*) but Landgrebe and Smith's reasons for doubting the computability of coping are independent of the worries raised by these other authors: rather, their reasons more closely resemble Hayek's reasons for doubting the possibility of effective central planning.⁹

Perhaps the closest we come overall to an argument that prefigures Landgrebe and Smith's is a suggestion in the final paragraph of Walsh (2017). Walsh observes that computational complexity may limit the possibilities for superintelligence. Even if the intellectual capacities of AI grow exponentially, this will not suffice to allow AI to run super-exponential algorithms (in less than super-exponential time), let alone de-

cide undecidable problems. We can construe Landgrebe and Smith as arguing that AI intelligence would indeed have to do something along those lines. To this argument we now turn.

3. THE LANDGREBE-SMITH ARGUMENT PRESENTED

In this section, I will present the key terms of the Landgrebe-Smith master arguments and then summarize their argument for premise (1). Here again are the master arguments:

The Coping Argument

1. There is a non-empty class C of complex systems that cannot be adequately and synoptically computationally modelled.
2. We cope.
3. If our brains are computers, we cope iff we adequately and synoptically computationally model some complex systems from C (those in our environmental niche).
4. Therefore, our brains are not computers. and then:

The Emulation Argument

1. There is a non-empty class C of complex systems that cannot be adequately and synoptically computationally modelled.
5. Our brains are systems in class C.
6. If our brains cannot be adequately and synoptically modelled, they cannot be emulated *in silico* with the requisite fidelity for whole brain emulation.
7. If our brains cannot be emulated *in silico* with the requisite fidelity for whole brain emulation, then our brains are not computers.
8. Therefore, our brains are not computers.

I will now say more about the notion of a complex (dynamical) system, the notion of ‘adequate and synoptic computational modelling’, and what is at stake in claiming that our brains are computers.

3.1 Complex Dynamical Systems

We mean ‘complex dynamical systems’ in the sense of complex systems theory. To begin, we can consider a three-body problem in physics. As readers familiar with Cixin Liu’s eponymous novels may be aware (spoiler alert), life on a planet in a three-star system would be very unpredictable.¹⁰

There is a specific mathematical reason for this. One relevant feature is that many systems of partial differential equations, such as those describing typical n-body problems, are not analytically solvable. In other words, no known mathematics allows us to determine a function expressed in terms of elementary operations from which we can actually calculate the exact trajectory for such a system from the differential equations specifying constraints on the rate of change of that trajectory, not even if we know the system’s exact initial conditions.

However, intractability rather than analytic unsolvability is the real problem. A system’s equations might be analytically solvable, but still hopelessly intractable. The more general and more pressing question is whether adequate approximations to a system’s equations are available. Indeed they often are, even where the equations in question are not analytically solvable, and the systems in question are potentially chaotic. NASA, for example, has recently managed to land one spacecraft, OSIRIS-REx, on an asteroid 200 million miles from earth, collect samples, and then deliver them back home, and to use another spacecraft, DART,

to successfully redirect the trajectory of an asteroid 160 meters in diameter, 7 million miles from earth, while travelling at 14,000 miles per hour.¹¹

But astronomical n-body systems are not as complex as it gets. Even NASA cannot consistently beat the stock market or predict the weather a month out. With an n-body system, we've fixed the elements of the system, and so we've fixed its phase space (and it is a comparatively low-dimensional one). But in the general case, complex systems are open, in constant interaction with their environments, undergoing continual change in which things are "in" the system and which are not. This can make the modelling even more intractable.

Landgrebe and Smith (in section 7.5.2) identify seven features that contribute to the complexity (and thus the mathematical intractability) of complex systems, with biological complex systems in general satisfying all seven. These are: 1) change and evolutionary character, 2) element-dependent interactions, 3) force overlay, 4) non-ergodic phase spaces, 5) drivenness, 6) context-dependence, 7) chaos.

Change and Evolutionary Character (pp. 126-128): complex systems evolve in various ways: the system's boundaries can shift, new elements come, old elements go. In many cases complex systems can undergo change in the types of elements they contain or interactions they participate in.

Element-Dependent Interactions (pp. 128-129): Complex systems typically have different kinds of functionally individuated elements, e.g. the different roles played by proteins, kinases and ATP in phosphorylation (contrasted with the way that mass and velocity are all you need to chart all of the interactions of a Newtonian system). Elements of a system can also change their functions over time.

Force Overlay (pp. 130-131): Complex systems typically involve interactions between all four of the basic physical interactions (EM, gravity, strong and weak). *Non-Ergodic/Complex Phase Spaces* (pp. 131-132): we cannot predict the trajectory of a complex system over its phase space by averaging over volumes of that phase space.

Drivenness (pp. 132-136): A driven system is a system that does not generally converge to equilibrium, because it has access to an external energy source, at least over some span of time.

Context-Dependence (p. 137): the interface between a complex system and its environment is constantly changing, e.g., which elements are part of the system vs. part of the environment, or what states the system can occupy.

Chaos (pp. 137-138): chaotic systems are unpredictable, because small differences in initial conditions may lead to large differences down the road.

When I speak of the relevant class C of complex systems, I mean complex systems satisfying all seven of these constraints.

Landgrebe and Smith also distinguish between complex systems, as per the above, and what they call *logic systems*. They first gloss the latter notion as follows (p. 62): a logic system is 'a system such as a simple device engineered in such a way that its behaviour can be predicted using the equations of physics and the rules of logic' (i.e. a system that can be adequately and synoptically modelled). They then (p. 122) give what appears to be intended as a strict definition. However, the definition they give does not rule out that three-body systems are logic systems, so I presume that they only mean to give necessary rather than sufficient conditions. As I understand them, we do best to just understand 'logic system' as 'system that is not too complex to be adequately and synoptically computationally modelled' or perhaps 'system that meets at most a few of the seven criteria for complex systems listed above'.

3.2 Adequate and Synoptic Computational Modelling

What then is it to be adequately and synoptically computationally modelled? A synoptic model is one that can be used "to engineer (or emulate the behavior of) a system or a system component". A model is "... adequate relative to some set of specified requirements if it can be used to engineer an artefact, or to create an emulation, that satisfies all the requirements of that set. An adequate model can be either almost exact, where any deviation from its predictions is so small that it is not measurable in experiments; ... or approxi-

mative, but where the deviation of the model from reality is irrelevant to the fulfilment of the requirements for the satisfaction of which the model was constructed.”¹² Such a model is computational when it can be encoded in a Turing-computable algorithm, i.e. emulated by some Turing machine.

3.3 The Brain as a Computer

I am using the phrase ‘the brain is a computer’ as a stand-in for the claim that human-level intelligence is explainable in terms of computations (analog or digital) that we or our brains implement. For Landgrebe and Smith, the key notion is that of human-level intelligence.

They view human or human-level intelligence as a combination of two forms of intelligence, *primal* and *objectifying*. Primal intelligence is characterized by Stern (1920): “the ability to adapt to new situations”. Objectifying intelligence is characterized by Gottfredson (1997): ‘A very general mental capability that, among other things, involves the ability to reason, plan, solve problems, think abstractly, comprehend complex ideas, learn quickly, and learn from experience’.¹³

Landgrebe and Smith suggest that while all animals exhibit primal intelligence, only humans display objectifying intelligence. The notion of ‘objectifying intelligence’ is loaded: is there really a singular capability that underwrites all of those things, that humans have and animals don’t?¹⁴ I’m not sure, but nothing I say below will hinge on this: we can think of objectifying intelligence as a cluster of niche-relative abilities if we prefer.

3.4 The Case for Premise One

In this subsection I consider Landgrebe and Smith’s case for premise one: the premise that there is a class C of complex systems that cannot be adequately and synoptically computationally modelled. The class C of complex systems is the class of systems that exhibit all seven of the features enumerated above (force overlay, chaos, etc.). For brevity I will call complex systems in class C “fully complex systems”.

Why think that fully complex systems cannot be adequately and synoptically computationally modelled? Landgrebe and Smith give us two arguments: first a concise deductive argument, and then a sustained argument by cases.

The deductive argument comes on page 116. Here they argue that no Turing-computable algorithm can adequately and synoptically model a fully complex system: “...this is because to be computable it must be isomorphic to an algorithm which can be expressed using the basic recursive Church functions. Each model consisting of a combination of these functions is always a model of a logic system, even if the latter is used to approximate a complex system. When such a model is executed on a computer, it becomes a process of the logic system which is the computer itself, which is realised in the way the binary logic of the microprocessors and the other components of the computer operate. But this system is not a subsystem of the complex system it approximates.”¹⁵

Setting aside a question about the notion of isomorphism in play,¹⁶ this argument boils down to the claim that a model *in* a logic system can only be an (adequate and synoptic) model *of* a logic system. But that trades on a conflation between different kinds of model: concrete models or mock-ups on the one hand, and mathematical or computational models on the other (see Weisberg (2013)). There is a sense in which every system is a concrete mock-up of itself (of perfect fidelity): what David Lewis might call a *Lagadonian* model. Granting that digital computational systems like standard desktop computers are logic systems if anything is, it follows that a process running an algorithm on a digital computer is, and thus is a Lagadonian model of, a process of a logic system. But as a mathematical or computational model, the relevant process may also model something else that is not a process of a logic system.¹⁷

Here is another way of seeing that the deductive argument proves too much. Think about all of the known ways that computational models can themselves be complex. Consider, for example, cellular automata like the game of life.¹⁸ Aren’t those examples of complex-enough systems that can be implemented

in silico? Crucially such a model may have a simple generation rule, but you still have to actually run the simulation to see what it will do. Further, some argue that everything physical is computable, defending, e.g. digital physics and the Bold Physical Church-Turing Thesis.¹⁹ This thesis may be wrong, but it is doubtful that we can refute it with an argument as quick as the paragraph I've reproduced above. Maybe if we beef up the definition of 'adequate and synoptic' so that an adequate and synoptic model of a system allows you to make precise-enough predictions *without* having to actually run a complete simulation, *then* we can conclude there are no such models for complex systems. Crucially though, this would be to allow that the relevant complex systems can themselves *be* computational models, undercutting much of what Landgrebe and Smith hope to show.²⁰

So I'm inclined to set aside the deductive argument: it may have heuristic value but it is not load-bearing. Fortunately, the bulk of chapter 8 is devoted to a case-by-case assessment of why there is no known mathematics or form of computational model (analog or digital) that suffices for a range of specific modelling tasks concerning fully complex systems. Begin with the limitation that the partial differential equations for many dynamical systems have no analytic solutions. Compound on this chaos, dynamic phase spaces, non-ergodicity and so on, and you have a strong case that the things most likely to figure in the job description of an adequate and synoptic model would be unachievable at least some of the time. Maybe in many cases (e.g. NASA's astronomical feats mentioned above) we can approximate to the needed degree of precision (which would contradict the letter of premise one), but there are no guarantees that we will be able to do so in every case (which still preserves its spirit).

Crucially, while it isn't obvious precisely what this implies about intractability or non-computability or undecidability, it is known that some questions related to the mathematics of complex dynamical systems are indeed undecidable or noncomputable.²¹ Thus it would be surprising if adequate and synoptic models were not strictly speaking non-computable in at least some cases, and it is plausible that finding or using them would be intractable (i.e., NP rather than P, or worse) in many others.

There is room for further discussion here. It is not clear that the undecidable questions are always or ever going to be the ones we'd need to settle to arrive at adequate and synoptic models of systems we care about (this is a contextual notion, remember); it is not clear which complexity class, in general, we'd be talking about; and it is not clear whether the limits applying to digital computers also apply to analog computers—but the limits on modelling fully complex systems are certainly real, and they flow from matters of principle rather than simply a need for more parameters in our model. So it seems reasonable to doubt that we humans get by in the world by adequately and synoptically modelling such systems in all cases. For these reasons I'll grant premise (1) for the sake of argument.

4. EQUIVOCATION IN THE COPING ARGUMENT

I turn my attention to a critique of the coping argument. I will argue that there is a critical equivocation on the meaning of 'coping', such that on one disambiguation, premise two (the claim that we cope) is false, while on the other disambiguation, premise three (the claim that computational coping requires adequate and synoptic modelling) is false.

What exactly is coping? A first distinction to draw is that between bodily coping and intellectual coping. Bodily coping is not our subject here: I am not defending the claim that it is computation all the way down (though I am not abandoning that claim either: it is just not our subject today).

A second distinction is between coping in a generic sense: getting on as well as we do, whatever that boils down to—and the more specific idea that we cope intelligently with fully complex systems. This is where equivocation sneaks in. Of course we cope as well as we do, whatever that means—but it does not follow that we are particularly accomplished at dealing intelligently with fully complex systems.

Certainly we are surrounded by such systems (and also made up of them). What does it mean, though, to say that we cope (with these systems)? Here is the ambiguity. Think of all of the things that we could do if indeed (counter-possibly) we had adequate and synoptical models of these systems. Then we could, in effect,

engineer them without limit: we could use our models to make predictions as accurate as needed, and then exploit that information where needed. That would give us an astonishing amount of control over the natural world. In the remainder of this section, I will argue that if this is what ‘coping’ means, then we do not cope and premise two is false (§4.1) while on the other hand if this is not what ‘coping’ means, then computational models of coping (that do not require adequate and synoptic modelling of fully complex systems) are still on the table, and premise three is false, or anyway undermotivated (§4.2).

4.1 The Argument from Human Frailty

Let’s take some examples: the weather, the stock market, social dynamics. There are no good computational models that allow us to predict whether it will be raining on your street at exactly 3 p.m. a month from now. But our solution to this is not some elaborate set of dynamic mechanisms that allows us to adjust in advance to whatever the weather will be, though we cannot compute it: we just muddle through, throw both suntan lotion and an umbrella in the bag, and hope for the best. Likewise, no computational models reliably beat the stock market, but most people who invest also fail to beat the market, and those who win presumably have their limits too. Or take social dynamics: it may be impossible to model exactly how a complex social event like, say, a presidential election will turn out. It is tempting to think of the winners as playing high-dimensional chess. But generally they, too, are just muddling along.

As Kahneman, Tversky and others have taught us, our cognitive capacities are suffused with heuristics and biases that generally work well, but get us into all sorts of trouble when they don’t.²² We operate under bounded resources, and typically must satisfice rather than optimize. While some argue that less-is-more,²³ few dispute that we often rely on very simple (and easily computable) heuristics and shortcuts, and that these run the risk of falling short when we find ourselves in overly challenging situations.

This doesn’t mean that we don’t cope: it is rather, a conception of how we cope. But it suggests that the way we cope is not by somehow accommodating every fluctuation of the fully complex systems around us: we just have our rules of thumb for how to manage, and they’ve worked so far, more or less. While it is true that the social world is complex, language is complex, markets are complex, and so on, we accidentally alienate friends and unwittingly elect con-men; we rely on dying metaphors and sound like broken records; we blow our savings on stock market speculation.

One might stress the various ways in which humanity as a species has thrived because we have found work-arounds to avoid the need to adequately and synoptically model fully complex systems. One example is our success at niche construction: that is, outsourcing cognitive tasks to our environment, making them less complex and more tractable.²⁴ Perhaps on an anthropological scale this means that as a species we can fully cope with fully complex systems, but at the same time this underscores how as individuals we must instead learn to muddle through despite our inability to expertly manage them.²⁵

I conclude that while there is of course a sense in which we cope—incomprehensibly well, from the chimpanzee’s point of view—the claim that we cope intelligently with fully complex systems must be taken with a great deal of salt. It is clear that we apply rules and statistics (i.e. intuitions) and we muddle through. It is highly doubtful that we master the fully complex systems around us, exerting the degree of control that we would if, indeed, we had adequate and synoptic models of those systems.

4.2 The Argument from the General Usefulness of Computation

Let us now turn to the other horn of my dilemma. If we want premise two to come out true, we must construe ‘coping’ to mean something more attainable: something less than exerting full mastery over the fully complex systems around us. But then, I will argue, premise three is false.

Premise three is the claim that if our brains are computers, we cope iff we adequately and synoptically model fully complex systems in our environment. Why accept this premise?

The right-to-left direction of premise three seems correct: if our brains are computers (or even if not), then if we adequately and synoptically model the relevant systems (quickly enough), we cope.

The trick is in the other direction: establishing that this is the only way to cope (or would be, if our brains were computers). Given the strong reading of ‘coping’, the one on which I’ve just argued we can’t cope, that follows: a computational system that has the abilities associated with a certain kind of modelling presumably does something equivalent to carrying out that modelling.

But in light of what I’ve just argued, we don’t cope in that strong sense: we aren’t playing high-dimensional chess all the time; we are generally just muddling through. The remaining question is whether there are computational models of the coping abilities we do have, models that do not require that we adequately and synoptically model fully complex systems.

And there are. Indeed, some of them are inspired by dynamical systems theory: many of the models it suggests can be implemented computationally. For example, the Hopfield net and other models exploiting attractor dynamics can be implemented within artificial neural networks.²⁶ Even van Gelder’s favorite example, the Watt Governor, would count by Landgrebe and Smith’s lights as a logic system rather than a fully complex system: it is a reliable, predictable device used in the engineering of actual engines!²⁷

Most importantly of all, computers are impressive. Think of all of the things that you can do on yours. Our driving concern is whether the brain *is* a computer. Put another way, we want to know whether those of our successes that seem due to our intelligence are successes *in virtue* of our brain’s ability to take on the form of a computer, or to approximate one well enough. Of course, every computer has bugs and glitches owing to features of its implementation: circuits leak or misfire, and neurons, too, might not do quite what the algorithm they implement dictates. The real question before us is whether, where intelligence is concerned, the rest is just noise, or whether the god of intelligence is hiding in those gaps.

It is hard to imagine any way to come closer to achieving maximally general primal and objectifying intelligence than by developing the capacity to compute a wide range of algorithms. There are indeed several theorems attesting to this: first, the classical Church-Turing thesis, which says that all generally recursive functions are computable on a Turing machine. In chapter seven, Landgrebe and Smith present this as a limiting constraint. But in reality, it is quite an achievement. Second, the universal approximation theorem, which says that a neural network (as a model of computation) can approximate any continuous function. Again, Landgrebe and Smith sometimes talk as though this is a limiting constraint, because it shows that curve fitting is *all* that neural networks can do, when in reality it is a remarkable achievement, equipping one with a toolkit of very general applicability.²⁸

Another way to put the point is in terms of control. On Landgrebe and Smith’s picture it is the intractable chaos that looks like (non-ergodic) noise from the computational perspective that ultimately explains how we do well. But to do well and survive for a long time it is not enough to be complex and chaotic: just ask a hurricane. It is not enough for there to be a storm beneath you; you have to be able to harness the wind. And in thinking about how we might have evolved ways to harness the chaos or near-chaos that churns beneath us, it is hard to find a better or more systematic answer than that our components found ways to function as coordinated computational nodes.²⁹

To conclude, let me reiterate. I am not arguing that computers will surpass all thresholds and benchmarks or eat free lunches. I am suggesting that it is very hard to identify concrete examples of intellectual abilities that we have that computation couldn’t enable at least as well as any alternative that we can articulate. Accordingly, on the understanding of ‘coping’ on which premise two comes out true, premise three comes out false, or at least unsupported.

5. THE CASE AGAINST PREMISES SIX AND SEVEN: THE ARGUMENT FROM SYNTHETIC MUSIC

Above, I have challenged premises two and three of the argument from coping: once you distinguish bodily coping from intellectual coping, I'm not sure what it means to say that we intellectually cope with fully complex systems. Though our computational methods may not amount to adequate and synoptic models of those systems, they may still be the best we can do.

Still, Landgrebe and Smith's second master argument, the argument from emulation, remains. Maybe our intellectual coping abilities do not exceed what (embodied) computers could do, but still, it remains the case that we ourselves are fully complex systems and, arguably, digital computational systems are not. Perhaps this is all we need to establish that our brains are not computers.

But care is required. Here, I will argue that there is an ambiguity in the notion of 'emulated *in silico* with the requisite fidelity for whole brain emulation' as it occurs in premises six and seven of the argument from emulation: are we operating at the level of types or at the level of tokens? If types, then premise six is false. If tokens, then premise seven is false.

Adequately and synoptically modelling a specific brain, just like adequately and synoptically modelling a specific hurricane, entails being able to predict precisely what that specific token entity is going to do. Adequate and synoptic modelling of specific hurricanes is intractable because, except for a few limiting cases, you cannot predict the path of a hurricane more than a few hours ahead with enough certainty to, e.g., assess precisely which houses will be hit and which won't be. Of course you can capture the *kinds* of things that hurricanes do and generate a model hurricane whose behavior is representative of the distribution, but that isn't enough for arbitrarily precise prediction of what an actual hurricane will do in the actual future.

For the sake of argument, I'll grant that the same is true of brains: that is, I'll grant premise five, that brains are fully complex systems. There certainly are questions here: do brains really satisfy force overlay, in the relevant sense? Do our brains really occupy a chaotic state (as opposed to a critical or near-critical state) during optimal cognition, e.g., when we are faced with hard choices between competing options, with little to go on? This is not something we can know *a priori*: one can model decision making using non-chaotic sampling. However, there is modelling work³⁰ and neuroimaging data³¹ that suggest that cortical dynamics may indeed be chaotic or weakly chaotic in some cases. See for discussion O'Byrne and Jerbi (2022).³²

So for the sake of argument I'll grant that in cases like that, for the same reasons that we cannot predict the exact paths of hurricanes, we cannot adequately and synoptically model individual brains: no matter how much you try to build a model of, e.g., me, there will be some limit on the reliability of the predictions you make about what I'll choose, at least when I am confronted with a hard choice between competing options with little to go on. In such a case, let's grant, my brain is a fully complex system, or close enough.

If this is true, there is an important respect in which premise six is also true. You then cannot build a model of me that would be guaranteed to mimic me in all relevant respects, e.g., to make all of the choices I would make. In other words, premise six is true if we understand 'emulation *in silico* with the requisite fidelity for whole brain emulation' as a claim at the token level, describing what it would take to build a perfect mimic of me, specifically.

But on this reading, premise seven is false: a computational model may be adequate as a model of the relevant type (*Human-Level Intelligence*), by generating a representative token of that type, even if it is not adequate (in Landgrebe and Smith's strict sense) as a replica of me, specifically.

On the other hand, if we take 'emulation *in silico* with the requisite fidelity for whole brain emulation' to mean, "adequacy as a representative token of the relevant type" then it may be that even though our brains cannot individually be adequately and synoptically modelled, still there can be simulated brains that have the requisite fidelity, i.e., that are representative tokens of the relevant type (*Human-Level Intelligence*). On this reading, premise six is false.

Note, crucially, that we do not have to say anything about dynamical systems complexity at the level of dendrites or ion concentrations to account for why token emulation of humans is intractable. Token emula-

tion of artificial neural networks is also intractable, unless you know the exact model parameters. If the link between a given two neurons were just a little stronger, that could tip the scales on a few critical decisions the organism might make, leading to a cascade of major differences in behavior down the road.³³ Thus, present considerations do not even force us to grant that the subtype *Strictly Human Intelligence* has a different behavioral distribution from *Human-Level Artificial Intelligence*.

Still, we might ask, what if it turns out that the subtype *Strictly Human Intelligence* does have a different behavioral distribution from *Human-Level Artificial Intelligence*? Can we somehow use this to show that there is no such thing as human-level artificial intelligence? No. Here, an analogy may help.

Consider the state of popular music in the 1980s. Synthesizers had broken onto the scene, like the SP-1200 (used by the Wu-Tang Clan and Pete Rock, among others). The SP-1200 is heavily quantized, meaning that you can't make synthetic drumbeats that capture all of the rhythmic micronuance of human drummers. Hence the mechanical, synthetic sound, and why record sampling remained vital.³⁴

Now suppose a skeptic comes along and says that it is impossible to make music with such crude instruments. Such a skeptic might say: "you can only make music with these tools if you can use them to perfectly replicate the sound of a human drummer." In answer, we can refuse the challenge: the music made with the SP-1200 really is music, even though it differs somewhat from that made by human drummers. The skeptic's task is to show us that the most salient conceptual kind here is one that excludes this newer more quantized kind of music. But for that, the skeptic has to show that the differences actually make the new stuff not only different, but non-musical. And I don't see how that would go.³⁵

The question about human-level intelligence (i.e., whether our brains are computers) is analogous to the question about what it takes to make *music*, not to the question of what it takes to perfectly replicate the music of a specific human musician, or even the music of human musicians in general. The moral is that we cannot presuppose that generating a representative token of the type *Human-Level Intelligence* requires generating a representative token of the subtype *Strictly Human Intelligence*. Accordingly, it does not follow from the fact that we cannot create a computational model that passes for human that the general features in virtue of which humans are intelligent cannot be instantiated *in silico*.

Of course, it also does not follow that those general features *can* be so instantiated. But this is where our previous discussion of the coping argument figures in: if it turned out that, to understand how human-level intelligence copes with the world, we had to appeal to non-computational mechanisms, then we could conclude that computational models of human-level intelligence must fail to adequately and synoptically model the relevant type (i.e., the type *Human-Level Intelligence*). But this has not been established. Our brains might still be computers, because it would only follow that they are not computers if we could show that you cannot emulate them (or what is salient about them) with requisite fidelity *at the type level*, and this has not been shown.

6. CONCLUSIONS

Much remains to be said. A first point is that if I am correct, then not only does it remain open that the brain is a computer, it also remains open that the future contains artificial general intelligence, or artificial super-intelligence. Even if there are abstract upper bounds on how well computational intelligence can cope with its environment, it doesn't follow that there aren't other computational intelligences that perform better than us on a wide-range of relevant power-seeking tasks (just ask the chimpanzees).

For all of that, there are many important morals to be drawn from Landgrebe and Smith (2022) with which I agree. Maybe the differences between biological humans and robotic or digital ones are essential, individuating differences, much as synthetic music and analogue music are fundamentally different kinds of music. If so, then perhaps we must revisit the question of the survivability of uploading. However, even if we grant that there is characteristic differences between biological and artificial kinds of intelligences it remains to be seen that an individual could not survive a change from one to the other. This calls for further metaphysical exploration.

There are also nearer term morals that even committed computationalists can draw from Landgrebe and Smith (many emerging in the final chapter of Landgrebe and Smith 2022). An awareness of the limitations of computational intelligence, and of the differences between biological and synthetic varieties, may help us navigate many of the pressing nearer and longer-term social challenges posed by the development of AI, and help us find a coherent perspective from which to harmonize efforts in dealing with the vast range of such issues from bias and misinformation in current models to risks of rogue AI down the road.³⁶

NOTES

- 1 Landgrebe and Smith (2022), pp.112-113
- 2 See for discussion Penrose (1989), chapter five, Eliasmith (2000), Dyson (2001), Piccinini and Bahar (2013).
- 3 van Gelder and Port (1995), Van Gelder (1995), see for criticism Eliasmith (1996).
- 4 Thelen et al. (2001).
- 5 Simmering and Perone (2012).
- 6 Ji et al. (2023).
- 7 Lucas (1961), Penrose (1989).
- 8 See for discussion Benacerraf (1967), Gaifman (2000).
- 9 Perhaps we may suggest, borrowing a phrase from Coffa, that Landgrebe and Smith (2022) is the book that Dreyfus would have written, had he been Hayek.
- 10 Of particular note: in chapter 17 of the first novel, Emperor Qin, following a suggestion of Von Neumann's, turns his entire nation into a large computer (c.f. Block 1978): unfortunately the predictions of this computer turn out to be unreliable (Liu and Liu 2014).
- 11 Thanks to David Anderson for suggesting this example.
- 12 Landgrebe and Smith (2022), pp.112-113.
- 13 *id.*, p.37.
- 14 For more on the debate here see Chollet (2017), Yudkowsky (2017).
- 15 Landgrebe and Smith (2022), p. 116
- 16 The thesis of Turing computability concerns functions rather than algorithms. If an algorithm computes a function that is Turing-computable then any universal Turing machine can *emulate* the algorithm to compute the function, but this does not entail an isomorphism at the architectural or mechanistic level. This opens up a loophole: maybe you grant that some algorithm suffices for (human-level) intelligence—if run on sufficiently complex, eg analog or neuromorphic hardware—but mere emulations of that algorithm on Turing machines do not. However, this distinction is less compelling in the case of intelligence than it is in the case of consciousness.
- 17 Thanks to Axel Constant for discussion.
- 18 Wolfram (2002).
- 19 Piccinini and Maley (2021).
- 20 Thanks to Jordan O'Byrne for conversation on this point.
- 21 See again Pour-El and Richards (1981), Pour-El (1999), and see also Hainry (2009), Grac,a and Zhong (2021).
- 22 Tversky and Kahneman (1974).
- 23 Gigerenzer and Brighton (2009).
- 24 See Brooks (1990); Constant et al. (2022, 2018); Henrich (2015).
- 25 For a discussion of another type of work-around, from modelling our own uncertainty in different models see Deane (2021).
- 26 See again Ji et al. (2023).
- 27 Crucially, also, we aren't here debating whether the mechanisms in your digestive or immune systems are computers: we are talking about intelligence strictly speaking. The computationalist can grant that your digestive and immune systems are (non-computational) dynamically coupled components of fully complex systems, without

allowing the same thing about the processes that explain your intelligence (especially what Landgrebe and Smith call your objectifying intelligence).

- 28 Of note here: in another paper, Landgrebe and Smith (2021b) argue that foundational models like ChatGPT fail to capture linguistic nuance, context, or syntactical rules. While these systems certainly have limits, it is remarkable how context-sensitive they are, thanks to attentional mechanisms (see Søgaard 2022), and how far they have come in the two years since Landgrebe and Smith (2021b) was written (see Piantadosi 2023).
- 29 For related (but more developed) arguments along these lines see Eliasmith (2000) and Piccinini and Maley (2021).
- 30 See van Vreeswijk and Sompolinsky (1996).
- 31 See Toker et al. (2022).
- 32 Thanks to Jordan O’Byrne for discussion.
- 33 Compare Godfrey-Smith (2016), who argues that neural-level explanations of cognition are inadequate.
- 34 See for discussion Roholt (2014), Charnas (2022).
- 35 Also worthy of note: contemporary synthetic beatmaking, e.g. using techniques pioneered by J-Dilla, has achieved what we might call superhuman performance. See again Charnas (2022).
- 36 My thanks to David Anderson, Jaan Aru, Axel Constant, George Deane, Janna Hastings, Jordan O’Byrne and Steve Peterson for comments. This research was supported by generous grants from the Good Ventures Foundation, and grants from FRQSC and SSHRC.

REFERENCES

- Benacerraf, P. 1967. God, the devil, and Gödel. *The Monist*, 51(1):9-32.
- Block, N. 1978. Troubles with functionalism. *Minnesota Studies in the Philosophy of Science*, 9:261-325.
- Brooks, R. A. 1990. Elephants don’t play chess. *Robotics and Autonomous Systems*, 6(1):3-15. Designing Autonomous Agents.
- Cantwell Smith, B. 2019. *The Promise of Artificial Intelligence: Reckoning and Judgement*. Cambridge, MA: MIT Press.
- Charnas, D. 2022. *Dilla Time: The Life and Afterlife of J Dilla, the Hip-Hop Producer Who Reinvented Rhythm*. New York: Farrar, Straus and Giroux.
- Chollet, F. 2017. The impossibility of intelligence explosion. *Medium*.
- Constant, A., Clark, A., Kirchhoff, M., and Friston, K. J. 2022. Extended active inference: Constructing predictive cognition beyond skulls. *Mind and Language*, 37(3):373-394.
- Constant, A., Ramstead, M., Veissière, S., Campbell, J., and Friston, K. 2018. A variational approach to niche construction. *Journals of the Royal Society Interface*, 15:1-14.
- Deane, G. 2021. Consciousness in active inference: Deep self-models, other minds, and the challenge of psychedelic-induced ego-dissolution. *Neuroscience of Consciousness*, 2021(2):niab024.
- Dreyfus, H. L. 1992. *What Computers Still Can’t Do: A Critique of Artificial Reason*. Cambridge, MA: MIT Press.
- Dyson, F. 2001. Is life analog or digital? *Edge*.
- Eliasmith, C. 1996. The third contender: A critical examination of the dynamicist theory of cognition. *Philosophical Psychology*, 9(4):441-63.
- _____. 2000. Is the brain analog or digital? *Cognitive Science Quarterly*, 1(2):147-170.
- Favela, L. H. 2020. Dynamical systems theory in cognitive science and neuroscience. *Philosophy Compass*, 15(8):e12695.
- Gaifman, H. 2000. What Gödel’s incompleteness result does and does not show. *Journal of Philosophy*, 97(8):462.
- Gerard, R. 1951. Some of the problems concerning digital notions in the central nervous system. In: von Foerster, H., ed. *The Stanford Encyclopedia of Philosophy*. New York, Macy Foundation, Summer 2021 edition.
- Gibson, J. J. 1979. *The Ecological Approach to Visual Perception: Classic Edition*. Boston: Houghton Mifflin.
- Gigerenzer, G. and Brighton, H. 2009. Homo heuristics: Why biased minds make better inferences. *Topics in Cognitive Science*, 1(1):107-143.
- Godfrey-Smith, P. 2016. Mind, matter, and metabolism. *Journal of Philosophy*, 113(10):481-506.
- Gottfredson, L. S. 1997. Mainstream science on intelligence. *Intelligence*, 24:13-23.
- Gra,ca, D. S. and Zhong, N. 2021. *Computability of Differential Equations*, pages 71-99. Cham: Springer.
- Hainry, E. 2009. Decidability and undecidability in dynamical systems. *HAL-Inria*.
- Henrich, J. 2015. *The Secret of Our Success: How Culture Is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*. Princeton: Princeton University Press.
- Ji, X., Elmoznino, E., Deane, G., Constant, A., Dumas, G., Lajoie, G., Simon, J., and Bengio, Y. 2023. Sources of richness and ineffability for phenomenally conscious states. *Neuroscience of Consciousness*, Vol. 2024(1). <https://doi.org/10.48550/arXiv.2302.06403>.
- Landgrebe, J. and Smith, B. 2021a. An argument for the impossibility of machine intelligence. *CoRR*, abs/2111.07765.

- _____. 2021b. Making AI meaningful again. *Synthese*, 198(March):2061-2081.
- _____. 2022. *Why Machines Will Never Rule the World: Artificial Intelligence Without Fear*. Abingdon: Routledge.
- Larson, E. 2021. *The myth of artificial intelligence: why computers can't think the way we do*. Cambridge, MA: Belknap.
- Liu, C. and Liu, K. 2014. *The Three-Body Problem*. New York: A Tom Doherty Associates Book. Tor Publishing Group.
- Lucas, J. R. 1961. Minds, machines and Gödel. *Philosophy*, 36(137):112-127.
- Maley, C. J. 2018. Toward analog neural computation. *Minds and Machines*, 28(1):77-91.
- O'Byrne, J. and Jerbi, K. 2022. How critical is brain criticality? *Trends in neurosciences*, 45(11):820-837.
- Penrose, R. 1989. *The Emperor's New Mind*. Oxford: Oxford University Press.
- Piantadosi, S. 2023. Modern language models refute Chomsky's approach to language. *Lingbuzz*.
- Piccinini, G. and Bahar, S. 2013. Neural computation and the computational theory of cognition. *Cognitive Science*, 37(3):453-488.
- Piccinini, G. and Maley, C. 2021. Computation in Physical Systems. In: Zalta, E. N., (ed.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2021 edition.
- Pour-El, M. B. 1999. Chapter 13 - the structure of computability in analysis and physical theory: An extension of church's thesis. In: Griffor, E. R. (ed.), *Handbook of Computability Theory*, volume 140 of *Studies in Logic and the Foundations of Mathematics*, pp. 449-471. Amsterdam: Elsevier.
- Pour-El, M. B. and Richards, I. 1974. Abstract computability and its relation to the general purpose analog computer (some connections between logic, differential equations and analog computers). *Transactions of the American Mathematical Society*, 199:1-28.
- _____. 1981. The wave equation with computable initial data such that its unique solution is not computable. *Advances in Mathematics*, 39(3):215-239.
- Roholt, T. C. 2014. *Groove: A Phenomenology of Rhythmic Nuance*. London: Bloomsbury Academic.
- Simmering, V. R. and Perone, S. 2012. Working memory capacity as a dynamic process. *Frontiers in Psychology*, 3.
- Søgaard, A. 2022. Understanding models understanding language. *Synthese*, 200(6):1-16.
- Stern, W. 1920. *Die Intelligenz der Kinder und Jugendlichen und Methoden ihrer Untersuchung*. Leipzig: Barth Press.
- Thelen, E., Schöner, G., Scheier, C., and Smith, L. B. 2001. The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral and Brain Sciences*, 24(1):1-34.
- Thompson, E. 2007. *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Cambridge, MA: Harvard University Press.
- Toker, D., Pappas, I., Lendner, J. D., Frohlich, J., Mateos, D. M., Muthukumaraswamy, S., Carhart-Harris, R., Paff, M., Vespa, P. M., Monti, M. M., Sommer, F. T., Knight, R. T., and D'Esposito, M. 2022. Consciousness is supported by near-critical slow cortical electro-dynamics. *Proceedings of the National Academy of Sciences*, 119(7):e2024455119.
- Tversky, A. and Kahneman, D. 1974. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124-1131.
- Van Gelder, T. 1995. What might cognition be, if not computation? *Journal of Philosophy*, 92(7):345-381.
- van Gelder, T. and Port, R. (eds.). 1995. *Mind as Motion: Explorations in the Dynamics of Cognition*. Cambridge, MA: MIT Press.
- van Vreeswijk, C. and Sompolinsky, H. 1996. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274(5293):1724-1726.
- Walsh, T. 2017. The singularity may never be near. *AI Magazine*, 38(3):58-62.
- Weisberg, M. 2013. *Simulation and Similarity: Using Models to Understand the World*. New York: Oxford University Press.
- Wolfram, S. 2002. *A New Kind of Science*. Champaign, Illinois: Wolfram Media.
- Yudkowsky, E. 2017. A reply to Francois Chollet on intelligence explosion. *Intelligence.org*.