

# How is there a Physics of Information? On characterising physical evolution as information processing

O. J. E. Maroney  
C. G. Timpson

Faculty of Philosophy, University of Oxford,  
Radcliffe Humanities, Radcliffe Observatory Quarter, Woodstock Road,  
Oxford, OX2 6GG  
owen.maroney@philosophy.ox.ac.uk  
christopher.timpson@bnc.ox.ac.uk

March 15, 2017

## Abstract

We have a conundrum. The physical basis of information is clearly a highly active research area. Yet the power of information theory comes precisely from separating it from the detailed problems of building physical systems to perform information processing tasks. Developments in quantum information over the last two decades seem to have undermined this separation, leading to suggestions that information is itself a physical entity and must be part of our physical theories, with resource-cost implications. We will consider a variety of ways in which physics seems to affect computation, but will ultimately argue to the contrary: rejecting the claims that information is physical provides a better basis for understanding the fertile relationship between information theory and physics. Instead, we will argue that the physical resource costs of information processing are to be understood through the need to consider physically embodied agents for whom information processing tasks are performed. Doing so sheds light on what it takes for something to be implementing a computational or information-processing task of a given kind

## Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
<b>2</b>	<b>Why is there a Physics of Information?</b>	<b>2</b>
2.1	The Challenge of Landauer’s Principle . . . . .	3
2.2	The Challenge of Quantum Computing . . . . .	4
2.3	The Challenge of Quantum Information . . . . .	4
2.4	The Challenge of Exotic Models of Computation . . . . .	5
2.4.1	Analogue computation . . . . .	5
2.4.2	<i>Objet trouvé</i> computation . . . . .	5
2.4.3	Niagra Falls computation . . . . .	6
2.5	The Challenge of Different Physics . . . . .	6
2.6	What is the Basis of the Resource Cost Claim? . . . . .	7
<b>3</b>	<b>The Ambiguity of Representation.</b>	<b>8</b>
<b>4</b>	<b>The Five Criteria</b>	<b>10</b>
4.1	Quantum Computing and DiVincenzo . . . . .	10
4.2	Landauer’s Principle and Maroney . . . . .	10
4.3	Unconventional Computation and Horsman et al. . . . .	11
4.4	The Final Five Criteria . . . . .	11

<b>5</b>	<b>Information and the Physically Embodied Agent</b>	<b>13</b>
<b>6</b>	<b>What makes a Someone?</b>	<b>15</b>
6.1	Boltzmann’s Laptop falls into a black hole . . . . .	15
6.2	Automata and other things . . . . .	16
6.3	Epistemic Communities of Agents . . . . .	17
<b>7</b>	<b>Conclusion</b>	<b>17</b>

# 1 Introduction

Information theory, whether considered in terms of communication or computational problems, is based upon analysing information processing tasks in the abstract, separated from the details of the physical substrate used to instantiate the task. However, in the last several decades there has been significant growth in research in a field called ‘The Physics of Information’, whose core contention is the importance of understanding that information processing is instantiated through physical systems. One of the slogans associated with this field is that ‘Information is Physical’ [1, 2, 3, 4]. As the late Rolf Landauer asserted:

Information is not a disembodied abstract entity; it is always tied to a physical representation. [2]

Although the exact sense in which ‘Information is Physical’ is intended to be understood is often left somewhat ambiguous (cf. [5]), many of the specific claims made in this field involve there being fundamental physical resource costs associated with performing information processing tasks, which costs should be understood as arising from the nature of the abstract information processing task itself.

In this chapter we will address some of the puzzles concerning these claims. We will begin by reviewing various ways in which physics and information interconnect, and the associated claims that certain information processing tasks have intrinsic physical costs. We will suggest that in order to make good upon these claims, the physics of information needs criteria for determining when a physical system instantiates a given information processing task, and that these criteria must be independent of either the model of information processing involved, or the physical theory describing the systems.

We propose a five-fold criterion which is clearly sufficient—and we argue is also necessary—to underpin resource-cost claims. These criteria must, furthermore, address the well-known problem of of the ambiguity of representation: when does a physical system represent a particular computation, and not some other computation? We argue that our criteria are able to resolve some of the key ambiguities, by adhering closely to an understanding of the physics of information based not on information being a physical thing, but rather based upon the physical embodiment of agents who are using physical systems to perform the information processing tasks in question.

# 2 Why is there a Physics of Information?

*Ab initio*, information theory seems based upon separating the analysis of information processing tasks from the physical systems that perform those tasks. Both the technical notions of computability and of information originally came about through a move away from ties to specific physical models.

The definition of a Turing Machine [6] made it possible to study the problem of what is computable, and with what difficulty, without being constrained to the particular details of physical systems needed to perform the computations. The fact that a Universal Turing Machine can efficiently (i.e., with at most a polynomial time increase) compute any algorithm that can be computed by any particular Turing Machine, made it possible to categorise the complexity of a given computational problem on the basis of how long it would take a Universal Turing Machine to solve the problem (while bearing in mind that open questions still exist regarding what precisely that categorisation is), and made it possible to group computational problems into distinct families based upon which problems could be solved by conversion (in polynomial time) into some other

problem with a known solution strategy. This categorisation does not seem to care how the Universal Turing Machine is implemented, whether through a network of logic gates or through other means such as cellular automata, for example.

Similarly, Shannon’s coding theorems [7] made it possible to study information-communication problems without being tied to the particular details of how a signal might be physically coded. Instead of looking at examples of specific communication schemes, or the content of the messages themselves, Shannon simply looked at the frequencies  $f_i$  with which signals occurred, and showed that the best possible compressibility and rate of transmission of the signals could be given as a function of those frequencies. The result was an operational definition of the information content of a source of signals that was independent of the details of the particular message or medium:

$$H = - \sum_i f_i \ln f_i. \quad (1)$$

In both cases it seemed that a universal resource cost could be derived for the performance of certain tasks, independent of the specific physical means by which the tasks might be instantiated. But we now wish to consider a number challenges to the idea that minimum resource costs associated with information processing can be wholly separated in this kind of way from the physical problems of building systems to perform the tasks. To anticipate: we do not wish to suggest that there are not fairly familiar things which might be said in response to various of the challenges we shall note; however, what we shall ultimately be after—and what we currently find lacking—is a unified and principled response to them, in contrast to the case-by-case treatments which one typically finds.

## 2.1 The Challenge of Landauer’s Principle

One of the key starting points in the move to consider information as physically embodied—or physically embroiled—in some significant manner comes from the thermal properties of computation. Computers typically give off large quantities of heat. The question of whether there is a non-trivial minimum quantity of heat that a computer must emit started to receive attention from the 1940’s [8, 9, 10, 11] before Landauer [12] proposed what is now generally regarded as being the correct way to understand this:

$$\Delta Q \geq -kT\Delta H. \quad (2)$$

The heat ( $\Delta Q$ ) generated by a computation—however implemented—is at least equal to the reduction in Shannon information over the course of the computation ( $\Delta H$ ) multiplied by the temperature of the environment in which the heat is generated. Given the similarity in form of the Shannon information measure and the Gibbs entropy measure

$$S = - \int dx dp P(x, p) \ln P(x, p), \quad (3)$$

it might seem entirely plausible that this should be the case. However, controversy has arisen [13, 14, 15, 16, 17, 18, 19, 20, 21, 22] over what exact assumptions could underly the claim than *any* physical system—of whatever kind—instantiating a particular computation must bear a thermodynamic resource cost.

Whilst it is fair to say that much of the controversy on this question has, at heart, concerned whether the kind of probabilities that are used in (1) are really of the same kind as are used in (3), there is also the question (more interesting for our current purposes) of how one can evaluate the generalisation ‘any physical system’. As Norton complains about statements of Landauer’s Principle:

One does not so much learn the general principle, as one gets the hang of the sorts of cases to which it can be applied. [16, p.383]

While there have been a number of attempts [23, 18, 24, 17] to address this complaint and prove Landauer’s Principle from general physical principles, there is a remarkable lack of agreement in such attempts as to what kind of assumptions can be taken as valid in constructing such a proof.

## 2.2 The Challenge of Quantum Computing

The next challenge comes from a concrete example of how a change in physics has led to a change in the minimum resource cost associated with an information processing task. The computational complexity of a given problem is, as we have mentioned, based on how quickly the best possible computer program, running on a Universal Turing Machine, will solve that problem as the size of the input to the problem increases. Simple cases, such as adding several numbers, increase in direct proportion to size and quantity of numbers being added. More complicated problems increase in time polynomially on the size of the input, and particularly complex problems can scale exponentially or even factorially with the size of the problem.

For a Universal Turing Machine, the problem of multiplying two numbers together increases polynomially as the size of the numbers increase. However, the problem of taking a number and finding if it factorises into two numbers increases super-polynomially.<sup>1</sup>

Then research in the 1980s and 1990s [25, 26, 27] discovered that a computer based upon the principles of quantum theory could factorise a number into its primes in only polynomial time (with an equivalent error to a Probabilistic Turing Machine). Constructing a physical system that is based upon primitive quantum operations, rather than classical logic gates, seems to lead to the possibility of an exponential increase in the speed by which some problems can be solved. To date, no known algorithm operating on a Universal or Probabilistic Turing Machine can perform this feat. The fundamental resource cost associated with solving the computational problem seems to be dependent upon physics.

But if a Universal (or Probabilistic) Turing Machine has hidden physical assumption, then what has happened to the idea of computational complexity as a feature which floats free of the physical substrate? And if it is the non-classical logic instantiated by quantum computers that allows the speed-up, how sure can we be that a classical physical system cannot be built which would allow efficient instantiation of equivalent non-classical logic gates?

## 2.3 The Challenge of Quantum Information

In tandem with the generalisation of the Universal Turing Machine to a Universal Quantum Computer, the second pillar of information theory—the Shannon Information measure—has also had to be revised in the face of quantum theory. This is, in some ways, even more remarkable. While no rigorous proof had ever been presented that a Universal Turing Machine really was efficiently equivalent to any other computational process (relying more on a plausibility argument that every realistic process had been shown to be so equivalent), in the case of Shannon Information, the assumptions seemed much weaker and the conclusion correspondingly stronger. The entire purpose of Shannon’s theorem is to characterise the signal capacity of a communication channel independently of the physical means by which the information is transmitted.

However, in the mid-1990s [28, 29] it was shown that a communication channel sending signals in quantum states could achieve a transmission rate that exceeded that of the Shannon theorem: if a quantum signal is represented by the quantum state  $\psi_i$ , then the quantum information capacity of a channel is given by

$$H_Q = -\text{tr}[\rho \ln \rho], \quad (4)$$

where  $\rho = \sum_i p_i \psi_i$ , and  $H_Q \leq H$ , with equality only occurring when the quantum signal states are all mutually orthogonal. For a given size of channel, that is, many more distinct quantum states could be sent intact from sender to receiver than would be possible classically.

The same problem is therefore raised as by the case of quantum computation: the Shannon information measure was intended to characterise a fundamental resource cost that was independent of the physical means by which the task was instantiated. But the change in physics, from classical to quantum, seems to change this fundamental resource cost.

---

<sup>1</sup>Note that for ease of comparison with the quantum case we should extend the notion of the Universal Turing Machine to a Probabilistic Turing Machine—which is equivalent to a Universal Turing Machine equipped with the ability to perform a fair coin toss. It is remarkable that there are problems that a Probabilistic Turing Machine is known to be able to solve, to within some error, exponentially faster than the best known algorithms on a Universal Turing Machine. However, factorising a number into its primes is—importantly—not one of those problems.

## 2.4 The Challenge of Exotic Models of Computation

While the challenges of the last two sections involve genuine advances in physics that have altered our understanding of the resource costs of performing information processing tasks, the next two sections will raise families of challenges which have been known, and—on a case by case basis—largely dismissed, for some time. We raise these challenges not to argue that they should not be discounted, but rather to start look for whether there is a *general* set of principles that explains why they should be disposed of. The first family of challenges is raised by the possibility that—before we even come to the question of the Universal Quantum Computer—the Universal Turing Machine might not, in fact, adequately characterise the best, or perhaps the only, means by which a physical system can instantiate a computation. We can consider various alternative *exotic models*. Are the claims that there are intrinsic resource costs such as Landauer’s Principle or computational complexity, just consequences of the limitations of particular models for how information processing is physically instantiated? Or can they be considered more general claims about any possible physical instantiation of information processing?

### 2.4.1 Analogue computation

The most well known challenge in this family is provided by analogue computation. Much of the proof of what is computable, and of the complexity of computation, for a Universal Turing Machine, rests upon the fact that there is only a countable infinity of data storage available to it. An analogue computer could potentially have a continuous infinity of information storage available, so these proofs would not apply. Analogue computers are generally discounted as realistic possibilities due to the effects of noise. In the presence of any finite amount of noise, the computational state of an analogue computer would be knocked into some different state, destroying the reliability of the computation. To be robust against such noise, the computational states would have to be represented by finite regions of the physical state space, and so the analogue computer would be reduced to a countably infinite amount of data storage. (More generally, any computational speed up due to analogue computing seems to require a resource cost of an exponentially increasing precision of the measurement of the computational state). Whilst this objection sounds *prima facie* plausible, there remains a concern about the basis on which the challenge is being answered: while we do not doubt the necessity of including noise in practical problems, the question might be raised as to why noise is, in principle, an unavoidable problem. The existence of noise is not written into the laws of classical physics, and quantum computing—for example—generally requires the existence of decoherence free subspaces to proceed: where is the proof that ‘noise-free subspaces’ are not possible in classical physics, allowing analogue computing to proceed?

### 2.4.2 *Objet trouvé* computation

A second challenge in this family comes from the suggestion that there might be physical systems that happen to be able to solve a complex problem in a short period of time. Suppose, for example, that it was possible to come across, or to build, a quantum system with an atom confined in a potential, where the eigenvalues of the Hamiltonian of the system correspond to the sequence of prime numbers [30]. Then if one wished to test whether an integer was prime or not—a computationally hard problem—one could simply shine a laser beam on the system, with a frequency that matches the integer in question. If the system resonates, then the integer is prime (interestingly, this would demonstrate that a number was not prime *without* determining its prime factors). And the time taken to perform this test would not seem to have any dependence on the size of the number being tested.

Two responses to this kind example can readily be discerned. The first would note that the precision of both the laser frequency and the potential would need to be very fine, to avoid a non-prime numbered frequency resonating a nearby prime number eigenvalue. This seems to be a similar kind of objection to the noise objection to analogue computing. The second is a concern as to whether it is genuinely possible to calculate the potential required accurately without already knowing the sequence of prime numbers. In other words, to build the system it might be necessary first already to have solved the problem itself, in which case solving the problem requires simply looking up the previously-found answer, rather than deploying a laser. (See [5, §6.4] for discussion of a related example of Nielsen’s [31].)

### 2.4.3 Niagra Falls computation

The next cases adapt a well known problem that has been raised in the literature on the philosophy of mind: the problem of whether it is even possible to state unambiguously whether a given physical system is uniquely instantiating a particular computation, and not some other [32, 33]. We will have more to say on this later (see Section 3). For now, let us start by considering Niagara Falls. Standing and looking at Niagara Falls, you will be observing a large amount of matter going through a great number of complex physical transformations. It is reasonable to suppose that, for any given computation you care to think of, somewhere in Niagara Falls there will be some part of the flow of water whose complex transformations match the abstract transformations required to perform the computation. Is Niagara Falls factorising large numbers into their primes, then?

The problem can be made more absurd by noting two further points: any classical computation can be efficiently performed by a reversible computer [34, 35] and any reversible computation corresponds to the action of a permutation on the computational states. Now, we can take a sufficiently large, but stationary system, and after labelling which physical states represent the input states for a computational problem, simply re-label the physical states as representing output states that correspond to the appropriate permutation of their input states. Call this *Stationary Computing*. Now the stationary system is apparently performing any computation by simply sitting there, being stationary. Such a conclusion is clearly absurd. What has gone wrong? Two considerations present themselves. The first is that we cannot read out the answer from the stationary computer unless we already know the relevant permutations, and this would correspond to us already having solved the relevant computational problem ourselves independently. Or second, if instead of this we were handed a code book that allowed us to decode the answer from the output state (or equivalently, encode the input state) then either the code is simply a look-up table that requires us to have already solved the problem as before, or else we will find that the decoding process will in fact be the computation, and all the resource costs have simply been transferred from the system to the decoding.

## 2.5 The Challenge of Different Physics

We now turn to the final set of challenges: quantum information and quantum computing appear to show that both the notions of Shannon Information and of Turing Computability, despite having originally been intended to be abstract analyses that floated free of physical assumptions, must have had hidden physical assumptions buried in them all along. If this can happen once, how are we sure that there are not further apparently innocuous physical assumptions we are making in our analyses of information and computation? And what happens if we take into account other physics, such as relativity, or even the possibility of our current best theories being replaced by entirely new physics?

One example in this line is that of hypercomputation. The notion of time resource associated with the hierarchy of complexity classes is not based upon the passage of physical time, but rather upon the number of elementary computational steps that are required. If the length of time per computational step is constant, then things are straightforward. But various hypercomputational models exploit the possibility that if each successive computational step takes a shorter and shorter length of time, then even a infinite number of computational steps might take place within a finite length of time. Is this physically possible? While quantum theory appears to provide ‘speed bumps’ [36], that would rule this out, general relativity seems to allow the possibility of space-times in which hypercomputation is possible [37, 38]. In a Malament–Hogarth space-time (which is not globally hyperbolic) one observer can, in a finite proper time, see the time line of a second system pass through an infinite proper time (before hitting a singularity). If the second system is a Universal Turing Machine, this can be used to solve the Halting Problem, which is formally unsolvable using a Universal Turing Machine. Here then, not only would complexity classes differ from the standard case, but so also would the computability class: the set of problems which can be solved at all.

If solutions to general relativity that are not globally hyperbolic are considered physically possible, then this also raises the possibility of computation involving closed-time like curves, which would again present a very different landscape [39, 40, 41, 42].

Finally, we note that even in the absence of experimental evidence, it is possible to reason about theories that might replace quantum theory in the same manner that quantum theory replaced

classical physics. Generalised Probabilistic Theories [43, 44] would allow information processing to take place differently to quantum theory. Once one has accepted that these things can change because quantum theory comes along, what is to stop one believing that things can change further if something else comes along? If, for example, a Popescu–Rohrlich non-local box [45] was discovered to be possible, would this have as dramatic an impact upon the definition of information as the discovery of quantum entanglement has done?

## 2.6 What is the Basis of the Resource Cost Claim?

To recap: the initial success of the theories of information and computation was that they seemed to provide a way of analysing information processing which was independent of particular models for how the processing task was to be physically instantiated. Resource costs could then be defined—such as channel requirements given data compression, the complexity of computation, or the thermal costs expressed in Landauer’s Principle—on the basis that any possible physical process that instantiates a given information processing task will have an intrinsic physical resource cost purely as a result of the properties of the (abstract) information processing task itself. At first blush, the advent of quantum information and quantum computation seemed an exciting new development within this picture; a sympathetic amendment or extension of it: here there were interesting new physical resource costs to incorporate into a *richer* physics of information. But equally—and catastrophically—one might see quantum information and computation instead as driving nails into the coffin of the physics of information. For once one has recognised that the details of physical instantiation cannot be ignored, then where do things end, as the challenges represented by our other examples illustrate? Even Landauer’s venerable principle might seem in question. How can any non-trivial statements about the resource costs of information processing be defined? How can there be a physics of information at all? Why isn’t there just physics?

It is dealing with the quantification ‘any possible physical process’ which is causing the trouble here. Any analysis that leads to a resource-cost claim might have made implicit physical assumptions about the underlying process: perhaps there exists some more clever method of instantiating the task which would avoid the cost. Thus analogue computing could do better, if it were possible to eliminate noise or efficiently measure states with exponential precision. Or to take another example, quantum computing uses a different family of primitive gates to classical computing, but the explanation for its doing better—its speed-up over the classical case—is not this detail of an abstract computing language, but rather the fact that discrete classical physical systems cannot efficiently instantiate the relevant quantum logic gates. Exploiting this fact gives the cleverer method of achieving the computational task.

The second aspect to the problem is that the notion of ‘possible physical process’ is itself a shifting terrain. Quantum theory has shown that abstract notions of information and computational complexity change if the fundamental physics changes. Why should we not expect general relativity, a quantum theory of gravity, or some post-quantum theory such as general probabilistic theories, similarly to force a fundamental revision in our notions of what kind of physical processes are possible for information processing? The very notion of information seems to have become dependent upon the physical theory used to model its processing. So how can we ask what the physical resource cost is of an information processing task, when the physics itself may be changing?

In our view what is needed is a perspective on all this that allows one to step a bit further back. To make a claim that an information processing task has an intrinsic resource cost, one will need:

- An abstract characterisation of the information processing task. What are the inputs, the outputs, and the logical relations between them? Who is being informed of what (and by whom)?
- Necessary (and sufficient) criteria for determining when a physical system instantiates a given information processing task.
- A proof, within any given physical theory, whether those criteria can be met, and if so, at what resource cost.

and we need to be able to apply these to any information processing logic, instantiated by any physical theory.

The challenges to the physics of information we have considered so far have primarily been pressing on the first and third of these requirements. We will now turn to the second requirement.

### 3 The Ambiguity of Representation.

The next challenge we must immediately confront is that, even in the most well understood case (classical information processing on classical system) the question of necessary and sufficient criteria for a physical system to instantiate a computation turns out to be a somewhat difficult and controversial problem.

Let us take the simplest suggestion: we define a logical gate as a map from a set of input states to a set of output states. In the case of the AND gate, for example, the map is:

$A$	$B$	$A \cdot B$
0	0	0
0	1	0
1	0	0
1	1	1

Now we design a physical system to instantiate that operation, in the simplest manner imaginable. We pick three physical system, each with a large number of physical states,  $x_1, x_2$  and  $x_3$ , and we associate a region of the physical state space of each with a logical state, so that if the physical state of the first system is in the region  $\{x_1\}_0$ , that represents  $A$  being in logical state 0, and if the physical state of the second system is in the region  $\{x_2\}_1$ , that represents  $B$  being in logical state 1 etc.

Now all that is required of our physical instantiation of the AND gate is to find some physical evolution of the system,  $f$ , that evolves<sup>2</sup> physical state  $x$  into physical state  $y = f(x)$ , and which ensures that whenever the first two systems initially lie in the joint region  $\{x_1\}_A \{x_2\}_B$  then at the end of the process the third system will always lie in the region  $\{x_3\}_{(A \cdot B)}$ . More generally we can state this formally as: if the set of physical states  $\{x\}_\alpha$  represent the overall logical state  $\alpha$  then the physical evolution  $y = f(x)$  instantiates the logical map  $\beta = M(\alpha)$ , if and only if,  $\forall \alpha, \beta$ ,  $y = f(x) \in \{x\}_{M(\beta)}$  whenever  $x \in \{x\}_\alpha$ .

Now, the ‘ $\forall \alpha, \beta$ ’ is important here, when calculating the resource cost of some operation. It is essential that the counterfactual *if some other input had been used, the correct output would still have been reached* holds, if the physical system is to truly be instantiating the operation.

Consider the following scenario (we owe the example to James Ladyman). A colleague suggests to you that they have a machine that can add two numbers much faster than a normal computer. To demonstrate they take you into a room with racks and racks of boxes, each box labelled with two numbers, and possessed of a button and a display screen. When you suggest two numbers to add, they pick up the box labelled by those two numbers and press the button. Immediately, the screen on the box lights up displaying the sum of those two numbers. When you suggest two different numbers, they do the same but with a correspondingly different box.

The problem here is that there is no reason to believe that any box is really performing the addition operation at all—for all we can tell, there might as well just be a light inside, with the pre-calculated answer already printed on the screen ready to show when the screen is back-lit.<sup>3</sup> We need one physical process that could have started with any valid input, and ends with the corresponding correct output in each case.<sup>4</sup>

So: we now happily build our physical instantiation of an AND gate, satisfying the mapping criterion stated above, and excitedly show it to a colleague. She looks blankly at us and says

<sup>2</sup>We are treating the evolution as deterministic in this simple example. Generalising this formulation to a fundamental stochastic dynamics is straightforward, but adds nothing to the argument.

<sup>3</sup>For illustrative purposes we have just framed the issue in epistemic terms: ‘no reason to believe...for all we can tell’. But the point is actually about what it takes for a device to be implementing some logical operation; this is a modal—a counterfactually loaded—notation.

<sup>4</sup>The room-full of numbered boxes would, perhaps, be a valid instantiation of addition, provided that all the pre-processing associated with deciding *which* box to select was included in the evaluation of the resource cost. Of course, even in this case, the presumably finite number of boxes involved would mean this approach would not scale to the addition of large numbers.

‘What are you talking about? That’s an OR gate!’. What has happened is that the physical states that we took to represent logical state 0, our colleague took to represent logical state 1, and vice versa. And under that different choice of representation, the exact same physical system is indeed representing an OR gate.

Now it might be tempting to try to establish that the physical system either *really is* instantiating an AND gate, or *really is* instantiating an OR gate, but we suggest that this would be a mistake. Suppose we now want to build a larger operation, out of AND and OR gates, then we could quite happily use exactly the same kind of physical process for each of the gates, but just wire up the connections to the inputs and outputs differently in the two cases. So it seems that there just is a fundamental ambiguity in which kind of gate the physical process instantiates.

Unfortunately, things do not stop there. Another colleague looks at the same system, and claims we have built an entirely different device altogether, one that performs a complex matrix multiplication. And when asked to explain how, it turns out that he has chosen a division of the physical states that is entirely different from ours, when representing the logical states. This is the problem that underlies the ‘Niagara Falls’ computation we saw in Section 2.4. If we are just using the criterion that the physical evolution must reproduce the logical operation, with complete freedom as to how the the physical states represent the logical states, any sufficiently large complicated system could be argued to instantiate any computation whatsoever. And reflecting on the case of Stationary Computation mentioned earlier, this would even apply to a system sitting there doing nothing, since *any* computation can be simulated just by permuting which unevolving physical states are to count as which logical states.

Of course there are ways to resist this seemingly pathological prospect of Stationary Computing, perhaps the most obvious being that we would seem to have to know the answer to the computation already in order to know how to label the output states correctly. But for now, we want to retain much of the ambiguity of representation, firstly because it may be of practical use to be able to use the same physical system to instantiate different operations, and secondly because we should be wary of ruling out some particular exotic computational model too hastily, in case our justification for ruling it out would also rule out another seemingly exotic model which, on development, might have more than a little going for it. Consider:

### Stationary Quantum Computing

- Any quantum operation can be implemented by unitary on a sufficiently large Hilbert space.
- Unitaries are just a change of basis.
- So any quantum computation can be performed by just preparing a state and then measuring in a rotated basis.

Stationary quantum computing is, occasionally, a useful way of studying the possibilities of quantum computation, but also, when the complexities of actually performing the measurement through a series of elementary operations, is taken into account, this model starts to become very close to *Measurement Based* Quantum Computing [46, 47, 48], which is one of the more viable models for practical quantum computing.

Now: the problem of the ambiguity of representation has been studied extensively, with particular importance for debates around computational theories of cognition (see e.g., [32, 33, 49, 50, 51, 52, 53, 54]). If it is thought that the mind is basically computational in nature, a matter of some interesting software running on some physical (presumably largely neurological) hardware, then it becomes rather important to be able to identify what software is running, and in virtue of what kinds of facts it is the case—if indeed it is the case—that some particular piece of computation is being performed on some occasion as opposed to some other one, or indeed none at all. We do not propose to address this large issue here, interesting and important as it is. For us this problem of the ambiguity of representation is embedded within and quite naturally at home in information theory because—as we said at the start—the whole point about the success of the Church/Turing notions of computability, and of the success of the Shannon Coding theorems, was that they analysed information processing abstracted from the essentially arbitrary ways in which the information might physically be represented and operated upon. And this was a good thing.

## 4 The Five Criteria

We have seen that some facet of ambiguity of representation may be a good thing. But if our only requirement for a process to count as computational or as information processing is simply having an evolution of the physical states which reproduces the mapping of the logical states, then this will leave too much ambiguity to be able to analyse the fundamental resource costs of information processing, with it being unclear that it amounts to any restriction at all. Fortunately, more restrictive criteria have been developed in various contexts, and we now turn to examine these.

### 4.1 Quantum Computing and DiVincenzo

In the 1990s, when the field of quantum computing was really beginning to develop, there were a great many proposals for which kind of physical systems might be viable for a scalable quantum computer. DiVincenzo (1997) [55] clarified the criteria which a system needed to meet to be a realistic possibility. His criteria were (paraphrasing):

1. A Hilbert space representation of the quantum states. The proposed systems must not simply be classical systems representing quantum states, but actually use a physical structure with a Hilbert space to represent the quantum states.
2. Initialisation of the physical system into a standard state. In order to get the computation going, the quantum system had to first be put into a ‘fiducial starting quantum state’.
3. Noise tolerance, especially with respect to decoherence affects. To perform quantum operations successfully it was essential to ensure that quantum coherence was not lost during the operations, so the system must have a high degree of isolation from the environment. This proved to be one of the biggest problems for early proposals until the development of quantum error correction codes and fault tolerance[56, 57, 58].
4. Ability to perform all the needed unitary operations. By analogy to a classical computer, a network constructed from a small set of elementary, universal logic gates (such as AND and NOT), a quantum computer needed to be a network constructed from small set of elementary, universal quantum operations (many 2 qbit quantum gates, such as the quantum controlled-NOT gate, satisfied this requirement). The development of more exotic models of quantum computation, such as quantum annealing or measurement-based quantum computation, led to this requirement being relaxed slightly: the evolution of a quantum computer had only to be unitarily equivalent to a network model of a universal quantum computer.
5. Ability to perform a strong projective measurement to read the output. While in a classical system, the end of the computation simply leaves the device holding the answer, in a quantum system it is necessary to perform a quantum measurement, meaning something rather different than simply a further unitary evolution, so that the result of the computation could be read out.

### 4.2 Landauer’s Principle and Maroney

As described earlier (Section 2.1), one of the complaints raised about the resource claim associated with Landauer’s Principle is that there was no proof offered that any physical system must be of the kind to which Landauer’s Principle applied. In attempting to clarify the basis of Landauer’s Principle, one of us put forward necessary criteria for a physical system to implement a classical logical operation (Maroney, Section 1.2 2005 [15], Sections II.D and II.E 2009 [18]). While the criteria varied slightly between the two papers, they shared a number of features in common (numbering follows the criteria as stated in 2005):

1. That there are distinct regions of the physical state space, which are used to represent the distinct logical states (2009 Section II.D 1-5)
2. When the physical system is placed in a state in one of the regions, it does not leave that region unless acted on by an operation. (2009, Section II.D 6).

3. It should be possible to reliably set the system to be in a physical state within one of the regions, even if it is not possible to set it into a precise physical state (This criterion was not explicitly mentioned in 2009)
4. It must be possible to determine the region of the physical state space in which the system lies (2009, only implicitly in the requirements of Section II.D 3-4)
5. There must be physical interactions that allow the logical operations to be performed (2009 Section II.E).

Beyond the second requirement, of the stability of the representation of the logical states, noise tolerance was not explicitly referred to for the operations themselves.

### 4.3 Unconventional Computation and Horsman et al.

In parallel to our own thinking on these matters, Horsman et al. [59] were independently seeking to address the question of when exotic computational models (which they refer to as ‘unconventional computation’) could be properly understood as performing a computation. While their analysis was expressed in different terms, exploring an analogy between confidence in unconventional computation and confidence in physical theories, the criteria they identify match the criteria we were developing:

1. There exists a representation relation between the input and output states of the physical system and the abstract states of the computational model
2. That this representation results in a commutative diagram when the evolution of the physical states, according to a theory of the physical system in which there was a high degree of confidence, is compared to the abstract computation in the model
3. That the physical state of the system at the start is encoded with the input logical state through an initialisation process.
4. That the logical output state is read off from the physical state of the system at the end, through a decoding process.
5. That the high degree of confidence in the physical theory used to predict the evolution of the system, is based upon it having been experimentally tested to within a good error bound.

### 4.4 The Final Five Criteria

The various criteria listed above display some variation, because each was developed to analyse a particular kind of problem. However, one can see clear similarity across these different cases:

#### 1. Representation of logical states by physical states

The physical system has to have the right kind of physical states available to represent all the required logical states. This representation will need faithfully to preserve the relevant properties of the logical states. In the case of classical computational states, for example, they need to be distinct, so the physical representations must correspond to distinguishable regions of the physical state space. Quantum computational states, by contrast, will need to be represented by physical states which include the phase relationships between the logical states.

The question of what the relevant structural properties are will, of course, vary between different kinds of logics of information processing, but provided the relevant properties are there in the physical representation, there is the full freedom to choose how to implement this which the ambiguity of representation allows.

#### 2. Initialisation of physical states

There must be a controllable physical processes which can reliably set the system to be in each of the initial representative states. The resource costs (such as time, space or work) that this initialisation process uses must be counted as part of the resource cost of the information

processing task. So all pre-processing must be included, otherwise it is possible simply to transfer the costs of performing the task to an initial encoding operation, such as in the case of the ‘Stationary Computing Model’, with an initial encoding that actually performs the permutation part of the computation.

### 3. Equivalence of evolution

The information processing task will involve a set of operations which map<sup>5</sup> input logical states to output logical states. The physical processes must evolve the representative physical states equivalently to the mapping of the logical states. It is important that the physical system does not just instantiate a particular mapping of input to output state. All the counterfactual claims must also be true of the system: if it had been presented with a representation of a different input state, the physical process would have still produced the representation of the correct corresponding output state. The resource costs (time, space, work) associated with this are part of the resource cost of the information processing task.

### 4. Readout of results

It is important to understand this not just as specifying the representation of the final logical states. It is the requirement that the final output of the information processing task must be in a form that is directly accessible to anyone who needs the information, in the sense not only that the output be of a *kind* which is readable, but it also be fixed exactly what logical output is represented.

In classical information processing, it might be assumed that the internal logical states can be read off from the physical states without any problems. However, for quantum computation, this is not necessarily the case. A quantum computation would generally involve non-orthogonal states, which cannot be distinguished, and cannot be measured without destroying phase coherence and disrupting the computation. Here, then, final measurement stage is needed at the end of the computation. For models of analogue computation, the readout stage also cannot be ignored: the need to measure the physical state with an exponentially increasing precision is of course one of the principal limitations which makes analogue computing unfeasible.

Once again, the resource cost (time, space, work) of this final stage must also be included in overall the resource cost of the information processing task. All post-processing must be included, otherwise we have the problem of the ‘Stationary Computing Model’, where the resource costs are simply shifted to a final decoding step.

### 5. Error tolerance

As the physical system is initialised, evolves and is measured, errors will inevitably occur. The errors depend upon both the kind of physical system being used and the kind of information processing task. Error tolerances is a major issue in the construction of a quantum computer, where environmentally induced decoherence can rapidly destroy a systems ability to perform quantum operations within quantum error correction codes. Error tolerance is also an issue for analog computing, and is generally considered to be the primary reason it is not feasible. Although error tolerance does not attract much attention in studies of classical computation, that is simply because the problem has been practically solved, and highly fault tolerant classical computers exist. That it can still be a problem becomes much clearer when computers overheat: in modern computing centres, the air-conditioning is one of the significant constraints upon computing power. In fact, error tolerance was understood to be a potential problem in the early development of classical computing, leading von Neumann to prove that classical computing could be made error tolerant (cf. Aaronson [60]).

Criteria 1 and 3 are just the same criteria that we studied in Section 3. Criteria 5 adds error tolerance, but this does not affect the problems that were raised by the ambiguity of representation. If our criteria are to resolve the problems we have encountered, it must be due to the addition of Criteria 2 and 4: the initialisation and readout stages.

---

<sup>5</sup>This includes stochastic maps, so that probabilistic operations and machines are included.

## 5 Information and the Physically Embodied Agent

It is not hard to maintain that our Criteria 1, 3, and 5 above are necessary conditions—but do not seem to be sufficient—to ensure that the information processing task is being performed. Our contention is that Criteria 2 and 4 are also necessary conditions, and that together 1–5 constitute a sufficient condition.

But that not everyone might agree is made clear by the following passage, from a paper by Ladyman, which is also concerned with analysing when a physical system is performing a computation:

In practice of course it is only possible to use a system as a computer if: (a) the relevant physical states are distinguishable by us (with our measurement devices); and (b) it is possible for us to put the system into a chosen initial state so as to compute the function in question for it. However, in principle some reason must be given as to why a system with respect to which these constraints are not met cannot be considered to be computing nonetheless. [61]

Given our argument so far, it seems there would be trouble if some principled reason could not be offered here, or trouble at least for the physics of information. For we have tried to argue that physical resource-cost claims cannot reasonably be made absent conditions like our full range 1–5 above, specifically including Criteria 2 and 4, which latter pair Ladyman seems to want to deny as essential. So much the worse for the physics of information it might be maintained! But we wish to preserve a domain for the physics of information—properly understood—and we think that a principled reason can be offered for why Criteria 2 and 4 should be conceived to be necessary conditions for computation to be happening at all, not just for its being useful.

To develop these thoughts we will need to clarify further what we understand by an information processing task.

We began our discussion with the general problem of trying to understand how information processing tasks could come with physical resource costs, purely by virtue of the identity of the information processing task itself. One obvious solution to this which we have not mentioned so far would be to propose that information itself should be conceived of as an independent physical property, such as energy or charge, or perhaps as some new kind of stuff, a new material or quasi-material substance, perhaps. This might make the idea that there is a physical cost to information processing seem perfectly natural. However, it would also seem to require there to be an unambiguous matter of physical fact that a given physical system is instantiating a specific information processing task, and that seems to be problematic.

Here we take an alternative view of the relationship between information and physics, drawing on the ideas developed in [5]. It is more than a little useful, in our view, to reflect that the term ‘information’ is an abstract noun, and correlatively to note that:

Very often, abstract nouns arise as nominalisations of various adjectival or verbal forms... their function may be explained in terms of the conceptually simpler adjectives or verbs from which they derive... ‘information’ is to be explained in terms of the verb ‘to inform’. Information... is what is provided when somebody is informed of something. [5, p.11]

We may often seem to use phrases such as ‘information is gathered’, ‘information was communicated’, ‘information will be stored’, as if information were some object or property that we can search for, pick up, move around and put somewhere. But there are also situations where it is not clear how that conception could make sense: for example, if some data is encrypted, and the encrypted data is stored separately from the encryption keys, the information seems only to be present in the correlations between the encrypted data and the encryption key. Destroy either, and the information appears lost (thereby also raising questions like ‘where did the information go?’)

We do well to step back and think about information not *as* a physical object or property, but rather as pertaining to a process or an activity for informing: it is ‘what is provided when someone is informed of something’. So in each case, rather than use the abstract noun, we could more fundamentally refer instead to the process of informing:

- ‘information is gathered’ vs. ‘an information gathering task is performed’
- ‘information is communicated’ vs. ‘an information communication task is performed’
- ‘information is stored’ vs. ‘an information storage task is performed’
- ‘information is processed’ vs. ‘an information processing task is performed’ etc.

On this conception, the different kinds of information processing tasks are different means by which someone is informed of something (and by someone). These tasks might be very complicated: the someone who is being informed may also be the someone who is doing the informing, for example if they are using a computer to perform a calculation; a communication task may involve multiple partners trying to communicate different things to each other, whilst concealing things from other people; some tasks, such as information storage, may be defined by the need indefinitely to defer the informing stage.

The important point we wish to draw from this for our purposes in this chapter, is that when seen in this way, an information processing task is defined by the question of who is being informed of what, and by whom, and the question of how that task may be physically instantiated is going to need to include a specification of how the informers and informees—as physically embodied systems themselves—interact with the information processing system. It is this which enforces and underwrites Criteria 2 and 4 as necessary requirements (*pace* Ladyman), in our view.

It is key that these *someones* are physically embodied agents. To inform, they must be able physically to start the process of initialisation, and if they cannot do this, the information processing task cannot begin. To be informed, the information must be presented to them in a form that they can directly perceive. The readout stage is required to do this, and until that has happened, the information processing task cannot be said to have completed. If either of those processes fail, the information processing task has not succeeded. Equally, if both those processes succeed, and the remaining criteria are also met, then the information processing task has certainly taken place. Thus we argue, on this view of information, the five criteria are individually necessary and jointly sufficient for a physical system to instantiate a well-defined information processing task.

We can start to see a little of how this will work by looking again at the humble example of the AND gate. The same physical process could be used to instantiate an OR gate. The difference between the two cases seemed to be how the physical states were chosen to represent the logical states, which seemed arbitrary. But what our criteria now add—to resolve this arbitrariness—is the requirement that when the system is actually being used to instantiate a particular information processing task then the initialisation and readout stages will fix exactly how the physical system is being used. Specifying those stages as an integral part of the task lifts the ambiguity of representation. The question of whether the physical process is an AND or an OR gate cannot be determined simply by looking at its physical evolution alone, but must involve a broader context. The question will be determined instead by how, in a given instance, the physical system is actually being used to inform, in virtue of its relation to *someones* as informers and informees.

Once the initialisation and readout stages are included as necessary components in any information processing task, the answers to many of the challenges we outlined earlier become unified. We have already seen how including the initialisation and readout naturally answers the problem of Stationary Computing, where the actual computation is shifted into a calculation of how physical states must be permuted to represent either the input or output logical states. It also answers the problem of the Niagara Falls computer: while there may be complex physical evolutions taking place within the waterfall, there is no-one who is initialising the state of the falling water, nor anyone reading out the supposed result<sup>6</sup>. No-one is being informed of any computation by such physical processes, so they are not information processing tasks.

It might be objected that some information processing tasks do not seem to lead to anyone being informed, such as information storage. Even in these cases, though, we would argue, the definition of a readout stage is a necessary component in the task. If the data is stored in an encrypted or encoded form, or is stored on the internal degrees of freedom of a physical device, such as a hard disc drive, the definition of the information storage task must include the specification of how that

---

<sup>6</sup>We may also note that it is unclear if the counterfactual condition on the equivalence of evolution holds in this model: if, somehow, someone had been able to ‘initialise’ a portion of the water to a different state, would the subsequent evolution still have produced the correct outcome?

data can be extracted again, whether through decoding the data, or through physically reading the data and presenting it to another system. If the coding scheme were to be lost or destroyed, the data would be lost too as it would be impossible to access it, and the data would no longer be meaningfully described as stored. Equally, if the technology needed physically to extract the data were to be damaged or lost, it would be impossible to access the data, and the storage task would have failed. For an information storage task to succeed, it is necessary that at some point in the future the data will still be accessible to someone, and the specification of how this readout stage can be achieved is still part of the overall information processing task.

We also wish to note that the need to include the physical process of initialisation and readout within the definition of the information processing task, leaves a great deal of freedom to exploit the ambiguity of representation in the way the task is instantiated in-between input and output. The definition of the information processing task must refer to the logical inputs and outputs as they are perceived by the physically embodied agents, but the equivalence of evolution criterion just refers to the overall mapping of these inputs and outputs, and does not need to specify a particular method by which it is achieved: there may be many different algorithms for achieving the same information processing task. So we need not identify the information processing task with any particular algorithm or computational model. It may even be the case that a given physical process for achieving the task could be interpreted in more than one way, in the same manner that the same physical process could be interpreted as an AND gate or an OR gate. Provided the initialisation and readout stages define the manner in which the physically embodied agents are informed, however, the instantiation of the information processing task itself will not suffer from this ambiguity.

## 6 What makes a Someone?

Our proposal is that the question of how information processing tasks can be physically instantiated cannot be separated from the question of the existence of physically embodied agents who are informing and being informed by the task. We have argued that this means that initialisation and readout stages must be included within the definition of the task, and that this also provides a principled answer to problems involving the ambiguities of representation when calculating the resource costs associated with information processing. However, the question still arises, what kind of physically embodied system is needed to fulfil the role of a *someone* to inform or be informed? There may be danger of a lurking circularity or a descent into vacuity here. And is it really true that information processing cannot take place without such a system or systems?

### 6.1 Boltzmann's Laptop falls into a black hole

Consider the following somewhat unlikely scenario. A large cloud of interstellar dust is going through thermal fluctuations when, by chance, these fluctuations lead to a chunk of matter coalescing into a physical object that happens to be in exactly the same state as the laptop on which this chapter is being written. The keyboard of this fluctuation-born laptop gets bounced around, leading to this very sentence appearing on its screen. And then the laptop falls into a nearby black hole and no-one ever sees it.

Now it might seem reasonable to say that that the fluctuation-born laptop did exactly the same computation and exactly the same information processing tasks as the laptop on which this paper was written. But it is a consequence of our point of view that this is not the case: no-one was being informed of anything, so no information processing task was taking place. It just happened to be a random thing that happened in space.

We might draw the parallel here to familiar questions about representation and in particular to Putnam's famous example of an ant crawling across the desert and accidentally inscribing the face of Winston Churchill on the sand [62, chpt.1]. Is the pattern in the sand actually a representation of Winston Churchill or was it just a random pattern that happens to resemble something else? It seems that it must only be a surprising random pattern for there was of course no intention on the ant's part to make a representation of Winston Churchill, or indeed a representation of any kind; and anyway, the ant had never come into any kind of appropriate causal contact with Churchill, or anyone who knew him, or knew of him.

However, there is something more that can be said about the fluctuation-born laptop. Had a passing spaceship noticed the laptop before it fell into the black hole, and scooped it up, then they would still have a perfectly good laptop which they could then use to perform computations, including writing about computations. The fact that it had come into existence by a random nonintentional process would not matter.

Yet there is a difference between having the capacity to perform a task and actually performing the task. While the laptop was floating through the dust cloud, randomly looking—if there were anyone to look—like it was typing up a chapter it can be considered to have shown the physical capacity to perform an information processing task, but it was not actually performing any such task. Only were the laptop to be scooped up and put to use would it indeed be using that capacity to inform the passing space travellers.

Similarly, the physical system instantiating the AND gate has the capacity to be used as an AND gate or an OR gate. Left to itself, it is doing neither. Only when it is placed in the context of being used by someone in an information processing task can it be said actually to be operating as one gate or another.

Information measures are often just ways of quantifying the capacity of systems to perform various information processing tasks. We can examine a physical system and analyse and quantify whether it has the capacity to perform a particular information processing task whether or not anyone actually uses the system to perform that task. But we would argue that it is only actually performing an information processing task if it is actually *being used* to perform such a task<sup>7</sup>.

## 6.2 Automata and other things

Still, we must ask: what does it take to be a *someone*? What about non-human animals, what about artefacts or various other kinds of things in general? What about automata? What about a thermostat informing the boiler of the temperature of the room? What about information stored in DNA being read during ontogenesis? These last two are very familiar examples of situations in which people often use the language of information, although it may seem unclear who or what would be supposed to be being informed, by our criteria. Given that this use of language seems to have pragmatic value at the very least, can we really be arguing that it is nevertheless wrong?

Does a thermostat inform the boiler of the temperature of the room? Are thermostats and boilers the kind of physical systems that can count as a *someone*? We might wish to define a *someone* in terms of the network of causal relationships they are part of, or in terms of agency—their capacities to interact with and respond states of the world external to them. The problem here, for us, is that trying to make those suggestions precise will often lead to a definition that is itself couched in information processing terms (for example, in the physics literature, physically embodied agents are often defined in terms of ‘information gathering and utilising systems’). This risks either introducing a circularity, or a return to the problem of the ambiguity of representation, but now applied one level up, to the entire system which includes the physically embodied *someones* being informed.

It seems there are two possible directions one could go in response to this. The first involves firmly digging in one’s heels, and saying that these general kinds of systems are not properly-speaking agents at all, and they are not informers or informees. In this case, while information-talk about such systems might be a useful metaphor or analogy to employ, we should not regard it as anything more than that: a useful metaphor or analogy.<sup>8</sup> A problem with this direction is that it may look too restrictive, especially when one reflects that systems like the thermostat may actually involve physical items—such as programmable microchips—which we otherwise would have described as information processors. Perhaps we could refer to the overall process, in this case the overall process of regulating the temperature of the room, and attempt to cast that in terms of information processing for us, but such a move can seem rather forced.

The second direction is more liberal, again looking at the overall process, but seeing that it can be analysed into separate parts, some of which can be cast in terms of an information processing

---

<sup>7</sup>That our criteria seem almost—but not quite—to turn the answer to the question into a tautology is, to us, an indication that we are along the right lines!

<sup>8</sup>It makes a difference whether one says metaphor or analogy. Metaphors are not generally considered to be candidates for literal truth and they derive their communicative purpose precisely from this. Analogies by contrast are, typically, candidates for literal truth, involving a claim of the form: A is to B as C is to D, cf. [63].

task. If that information processing task would be the same if the systems informing and being informed could be replaced by physical systems (such as ourselves) that are uncontroversially identifiable as agents, without affecting the overall process, then it is valid to talk of that part of the process as being an information processing task. So, a thermostat may not be informing a boiler of something, but a boiler might be seen as using a thermostat as an information processing device, to inform itself about the temperature of the room, in just the same way that we might look at a thermometer to inform ourselves and decide whether or not to turn the heating on. The worry with this strategy is that too liberal an approach risks the possibility of everything and anything being cast in information processing terms, which obscures or obviates the point of the description.

### 6.3 Epistemic Communities of Agents

One of the key features of our physically embodied agents is that they must be able to initialise the process, and receive the readout. It is their physical attributes and abilities that ultimately fix the start and end points of the process. Does this mean that any physical resource cost associated with an information processing task will be different for different *someones*?

In practice, initialisation and readout procedures will be relatively insensitive to many—perhaps the majority—of the details of the specific physical agent being informed, as such agents will tend to form what we may term (with a nod to van Fraassen) *shared epistemic communities*. A shared epistemic community means a group of people, or entities, or agents, that can share and access the same information. They therefore have the ability to interact with the same information processing devices through the same kind of physical means. Put another way, we can think in terms of initialisation and readout being defined relative to a set of physically well-defined observers that share roughly the same kind of capacities to respond to and change the world surrounding them. It will be an objective physical fact about what the initialisation and readout capacities of a given epistemic community is.

But we can go further. In terms of physical abilities, humans, sheep and pigs (to take a random selection) do not differ so very greatly in terms of perception. We may not be able to communicate well with pigs and sheep, but insofar as we know of information processing tasks relevant to pigs and sheep, there will be a shared character of physical interaction with the information processing devices required. The resource costs of initialisation and readout will not be very sensitive to such details, and we would expect that they could be efficiently translated between the needs of these different epistemic communities.

What if we move further afield? What if we encounter aliens, with radically different perceptual capacities? First let us note: if the perceptions of the aliens depend upon a different physics to the physics of our perceptions, then we are just talking about the cost of information processing evaluated by a different physical theory. So let us assume that the aliens' perceptions come from the same physics, but using radically different physical processes to ours.

In this case it should be possible to analyse and understand those physical processes. We therefore conjecture: if aliens could perceive something, a talented experimenter can build initialisation and readout systems to use those some perceptual mechanisms, for us (and vice versa). So, assuming we had not unnecessarily restricted ourselves in some way when analysing fundamental resources costs, then a different shared epistemic community, based on alien perceptions but using the same underlying physics, should still yield the same fundamental physical resource costs.

## 7 Conclusion

So where does this leave our original problem? How can there be a physics of information whilst still holding on to the idea that information theory succeeds by abstracting the definition of information processing tasks away from the details of the physical substrates used to instantiate the tasks?

In this chapter we have approached the physics of information from the perspective of establishing physical resource costs associated with performing information processing tasks. To try and quantify these costs, we define an information processing task in terms of 'someone being informed of something by someone'. While acknowledging that there is a great deal of complexity that can be hidden in those terms, the *someones* involved are themselves physically embodied agents, and the physical means by which they start and end the process of informing must be included

within the specification of how the information processing task is to be instantiated. These are the initialisation and readout stages of information processing, which provide physical connections between the physically embodied agents and the information processing devices.

The physical means by which the task is performed does not need to be specified at this point: indeed, if the objective is to derive a fundamental resource cost, then specific models should be avoided. However, a proof of a physical resource cost cannot be wholly independent of the fundamental physics used to evaluate that cost. So any resource cost claim has to be made relative to a theory of physics:

- Many of the traditional resource-cost claims are based on the assumption of: classical physics, non-relativistic space-times, and unavoidable noise. Although physical systems for instantiating classical information processing devices are based on quantum theory (semi-conductor physics), the specifically quantum nature of the systems was not considered relevant for how the information processing was analysed.
- Quantum computing has generally been analysed on the assumption of: quantum physics, non-relativistic space-times, and models of decoherence that allow error correction. Recently, the problems of relativistic quantum information have been raised.
- Hypercomputational models such as Malament–Hogarth have been analysed in the context of classical, general relativistic space-times.
- Quantum gravity or general probabilistic theories would further change what is possible.

Once the relevant physical theory has been selected, the evaluation of the resource costs becomes the question of whether the five criteria can be met for the information processing task given that physical theory. This means, for example, models such as analogue computing or hypercomputation may seem viable when evaluated with respect to one kind of physical theory, but not with respect to another.

The greatest practical importance should obviously be attached to the assumptions that correspond to our current best physical theories, but the ever present possibility that these theories may need revision should make us cautious about ruling out exotic models of computation in principle. There is an interesting corollary here, though: actually building a system to perform some computational task with a resource cost below a minimum threshold could be taken as *prima facie* experimental evidence against a physical theory that sets that threshold.

Relative to a given physical theory, there may be some particular tasks that can be efficiently implemented, and may be efficiently combined as building blocks out of which all other tasks can be constructed. These building blocks may be used to define a logic of information processing that is well suited to the physical theory: thus Boolean logic gates are a basis for classical information processing, and quantum operations are a basis for quantum information processing. Equally, some resource costs may turn out to be particularly useful or ubiquitous within a given physical theory, and, like the Shannon and Schumacher measures, become regarded as a measure of information within that theory.

It should be noted that the consistency of a given abstract logic of information processing is not dependent upon physics: the only question is whether that logic can be efficiently implemented within a given physical theory. So information processing can continue to be studied in the abstract, independent of the means of physical instantiation. However, the interplay with physics means that sometimes physics will suggest new ways of analysing information processing tasks, enriching both fields, and logics which best match our current best physics will be the ones best suited for understanding the physical resource costs of real information processing devices.

## Acknowledgements

We thank the editors for the invitation to contribute to this volume, and for their patience during the process. This work was supported by a grant from the Templeton World Charity Foundation.

## References

- [1] Rolf Landauer. Information is physical. *Physics Today*, May:23–29, 1991.
- [2] Rolf Landauer. The physical nature of information. *Physics Letters A*, 217:188–193, 1996.
- [3] Seth Lloyd. *Programming the Universe*. Vintage Books, 2006.
- [4] Vlatko Vedral. *Decoding Reality: The Universe as Quantum Information*. Oxford University Press, 2010.
- [5] Christopher G Timpson. *Quantum Information Theory and the Foundations of Quantum Mechanics*. Oxford University Press, 2013.
- [6] Alan Turing. On Computable Numbers, with an application to the *Entscheidungsproblem*. *Proceedings of the London Mathematical Society*, 42:230–265, 1936.
- [7] Claude E Shannon. The mathematical theory of communication. *Bell Systems Technical Journal*, 27:379–423, 623–656, 1948.
- [8] John von Neumann. The role of high and of extremely high complication. In A.W. Burks, editor, *Theory of Self-Reproducing Automata*, pages 64–73. University of Illinois Press, 1966. Lecture originally delivered in 1949.
- [9] Léon Brillouin. Maxwell’s demon cannot operate: Information and entropy I. *Journal of Applied Physics*, 22:334–337, 1951.
- [10] J Rothstein. Information, measurement and quantum mechanics. *Science*, 114:171–175, 1951.
- [11] Dennis Gabor. Light and information. *Progress in Optics*, 1:111–153, 1964.
- [12] Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5:183–191, 1961.
- [13] John Earman and John D Norton. Exorcist XIV: The wrath of Maxwell’s Demon. Part I. From Maxwell to Szilard. *Studies in History and Philosophy of Modern Physics*, 29(4):435–471, 1998.
- [14] John Earman and John D Norton. Exorcist XIV: The wrath of Maxwell’s Demon. Part II. From Szilard to Landauer and beyond. *Studies in History and Philosophy of Modern Physics*, 30(1):1–40, 1999.
- [15] Owen J E Maroney. The (absence of a) relationship between thermodynamic and logical reversibility. *Studies in History and Philosophy of Modern Physics*, 36:355–374, 2005.
- [16] John D Norton. Eaters of the lotus: Landauer’s principle and the return of Maxwell’s demon. *Studies in History and Philosophy of Modern Physics*, 36:375–411, 2005.
- [17] James Ladyman, Stuart Presnell, Anthony J Short, and Berry Groisman. The connection between logical and thermodynamic irreversibility. *Studies in History and Philosophy of Modern Physics*, 38(1):58–79, 2006.
- [18] Owen J E Maroney. Generalizing Landauer’s principle. *Physical Review E*, 79, 2009.
- [19] John D Norton. Waiting for Landauer. *Studies in History and Philosophy of Modern Physics*, 42:184–198, 2011.
- [20] James Ladyman and Katie Robertson. Landauer defended: Reply to Norton. *Studies in History and Philosophy of Modern Physics*, 44:263–271, 2013.
- [21] John D Norton. The end of the thermodynamics of computation: A no go result. *Philosophy of Science*, 80:1182–1192, 2013.
- [22] James Ladyman and Katie Robertson. Going round in circles: Landauer vs. Norton on the thermodynamics of computation. *Entropy*, 16:2278–2290, 2014.

- [23] Barbara Piechocinska. Information erasure. *Physical Review A*, 61:1–9, 2000.
- [24] S Turgut. Relations between entropies produced in non-deterministic thermodynamic processes. *Physical Review E*, 79:041102, 2009.
- [25] Richard P Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6/7):467–488, 1982.
- [26] David Deutsch. Quantum theory, the Church-Turing Principle and the universal quantum computer. *Proceedings of the Royal Society of London A*, 400:97–117, 1985.
- [27] Peter W Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *Proceedings of the 35th Annual IEEE Symposium on Foundations of Computer Science*, 1994. See also arXiv:quant-ph/9508027.
- [28] Benjamin Schumacher. Quantum coding. *Physical Review A*, 51(4):2738, 1995.
- [29] Richard Jozsa and Benjamin Schumacher. A new proof of the quantum noiseless coding theorem. *Journal of Modern Optics*, 41:2343–2349, 1994.
- [30] G Mussardo. The quantum mechanical potential for the prime numbers, 1997. arxiv:cond-mat/9712010.
- [31] Michael A Nielsen. Computable functions, quantum measurements, and quantum dynamics. *Physical Review Letters*, 79(15):2915–2918, 1997.
- [32] Hilary Putnam. *Representation and Reality*. MIT Press, 1988.
- [33] John R Searle. *The Rediscovery of the Mind*. MIT Press, 1992.
- [34] Charles H Bennett. The thermodynamics of computation—A review. *International Journal of Theoretical Physics*, 12:905–940, 1982.
- [35] Edward Fredkin and Tommaso Toffoli. Conservative logic. *International Journal of Theoretical Physics*, 21(3/4), 1982.
- [36] Seth Lloyd. Ultimate physical limits to computation. *Nature*, 406:1047–1054, 2000.
- [37] John Earman and John Norton. Forever is a day: Supertasks in Pitowsky and Malament-Hogarth spacetimes. *Philosophy of Science*, 60:22–42, 1993.
- [38] Mark Hogarth. Non-Turing computers and non-Turing computability. *Philosophy of Science Supplementary*, I:126–138, 1994.
- [39] David Deutsch. Quantum mechanics near closed timelike lines. *Phys. Rev. D*, 44:3197–3217, 1991.
- [40] Todd A. Brun. Computers with closed timelike curves can solve hard problems efficiently. *Foundations of Physics Letters*, 16(3):245–253, 2003.
- [41] Dave Bacon. Quantum computational complexity in the presence of closed timelike curves. *Phys. Rev. A*, 70:032309, 2004.
- [42] Scott Aaronson and John Watrous. Closed timelike curves make quantum and classical computing equivalent. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 465(2102):631–647, 2009.
- [43] Jonathan S Barrett. Information processing in generalized probabilistic theories. *Physical Review A*, 75:032304, 2007.
- [44] H Barnum, J Barrett, L Orloff Clark, M Leifer, R Spekkens, N Stepanik, A Wilce, and R Wilke. Entropy and information causality in general probabilistic theories. *New Journal of Physics*, 12:033024, 2010.

- [45] Sandu Popescu and Daniel Rohrlich. Quantum nonlocality as an axiom. *Foundations of Physics*, 24:379–385, 1994.
- [46] Robert Raussendorf and Hans J Briegel. A one-way computer. *Physical Review Letters*, 86(22):5188–5191, 2001.
- [47] Robert Raussendorf, Daniel E Browne, and Hans J Briegel. Measurement-based quantum computation on cluster states. *Physical Review A*, 68:022312, 2003.
- [48] Simon C Benjamin, Brendon W Lovett, and J M Smith. Prospects for measurement-based quantum computing with solid state spins. *Laser and Photonics Reviews*, 3:556, 2009.
- [49] D J Chalmers. Does a rock implement every finite state automaton? *Synthese*, 108:309–333, 1996.
- [50] B Jack Copeland. What is computation? *Synthese*, 108:335–359, 1996.
- [51] O Shagrir. What is computer science about? *The Monist*, 82(1):131–149, 1999.
- [52] Mark D Sprevak. *Mind and World: A Realist Account of Computation in Cognitive Science*. PhD thesis, Cambridge University, 2005.
- [53] G Piccinini. Computation without representation. *Philosophical Studies*, 137:205–241, 2008.
- [54] Mark D Sprevak. Computation, individuation and the received view on representation. *Studies in the History and Philosophy of Science*, 41(3):260–270, 2010.
- [55] D. P. Divincenzo. Topics in quantum computers. In Lydia L. Sohn, Leo P. Kouwenhoven, and Gerd Schön, editors, *Mesoscopic Electron Transport*, pages 657–677. Springer Netherlands, Dordrecht, 1997.
- [56] Peter W. Shor. Scheme for reducing decoherence in quantum computer memory. *Physical Review A*, 52:R2493, 1995.
- [57] Andrew Steane. Multiple-particle interference and quantum error correction. *Proceedings of the Royal Society of London A*, 452(1954):2551–2577, 1996.
- [58] Dorit Aharonov and Michael Ben-Or. Fault-tolerant quantum computation with constant error rate. *SIAM Journal on Computing*, 38(4):1207–1282, 2008.
- [59] Clare Horsman, Susan Stepney, Rob C. Wagner, and Viv Kendon. When does a physical system compute? *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 470(2169), 2014.
- [60] Scott Aaronson. *Quantum Computing since Democritus*. Cambridge University Press, 2013.
- [61] J Ladyman. What does it mean to say that a physical system implements a computation? *Theoretical Computer Science*, 410(4–5):376, 2008.
- [62] Hilary Putnam. *Reason, Truth and History*. Cambridge University Press, 1982.
- [63] Roger M White. *Talking About God: Analogy and the Problem of Religious Language*. Ashgate, 2010.