



Deposited via The University of York.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/147382/>

Version: Accepted Version

Conference or Workshop Item:

Stepney, Susan and Kendon, Viv (2019) The role of the representational entity in physical computing. In: UCNC 2019, Tokyo, Japan, June 2019, 03-07 Jun 2019.

https://doi.org/10.1007/978-3-030-19311-9_18

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

The role of the representational entity in physical computing

Susan Stepney¹ and Viv Kendon²

¹ Department of Computer Science, University of York, UK

² Department of Physics, Durham University, UK

Abstract. We have developed abstraction/representation (AR) theory to answer the question “When does a physical system compute?” AR theory requires the existence of a *representational entity* (RE), but the vanilla theory does not explicitly include the RE in its definition of physical computing. Here we extend the theory by showing how the RE forms a linked complementary model to the physical computing model, and demonstrate its use in the case of intrinsic computing in a non-human RE: a bacterium.

1 Introduction

Many and diverse physical substrates are proposed for unconventional computing, from relativistic and quantum systems to chemical reactions and slime moulds, from carbon nanotubes to non-linear optical reservoir systems, from amorphous substrates to highly engineered devices, from general purpose analogue computers to one-shot devices. In another domain, biological systems are often said to perform information processing. In all these cases it is crucial to be able to determine when such substrates and systems are specifically computing, as opposed to merely undergoing the physical processes of that substrate.

In order to address this question, we have been developing *abstraction/representation* theory (AR theory). This is a framework in which science, engineering/technology, computing, and communication/signalling are all defined as a form of *representational activity*, requiring the fundamental use of the representation relation linking physical system and abstract model in order to define their operation [3, 7]. Within this framework, it is possible to distinguish scientific experimentation on a novel substrate, from the performance of computation by that substrate.

In work following on from the original definitions, [6] provides a high level overview, [4] delves into more philosophical aspects, and [5] presents an example of intrinsic computation: signalling in bacteria. Also see [3, 4] for references to the wider unconventional computing and philosophical literature.

AR theory requires the existence of a representational entity (RE) to support the representation relation. One issue glossed over in our previous descriptions of AR theory that becomes crucial when analysing computation in systems where the RE is not a human or conscious user, is the relationship between the physical

RE and the physical computer. Here we enrich AR theory by incorporating the RE explicitly, and showing how it relates to the physical computing process.

The structure of the paper is as follows. In §2 we summarise the current formulation of AR theory. In §3 we extend the theory to include the RE explicitly. In §4 we demonstrate how the extended theory allows us to capture and model intrinsic computing in a bacterium.

2 AR theory in a nutshell

2.1 Our view of physical computing

AR theory has been developed to answer the specific question of when a physical system is computing [3]. Its answer hinges on the relationship between an abstract object (a computation) and a physical object (a computer). It employs a language of relations, not from mathematical objects to mathematical objects (as is usual in mathematics and theoretical computer science), but between physical objects and those in the abstract domain. The core of AR theory is the *representation relation*, mapping from physical objects to abstract objects. Experimental science, engineering, and computing all require the interplay of abstract and physical objects via representation in such a way that their descriptive diagrams *commute* such that the same result can be gained through either physical or abstract evolutions (see §2.3). From this, Horsman et al [3] define computing as *the use of a physical system to predict the outcome of an abstract evolution*.

2.2 Representation

AR theory has physical objects in the domain of material systems, abstract objects (including mathematical and logical entities), and the representation relation that mediates between the two. The distinction between the two spaces, abstract and physical, is fundamental in the theory, as is their connection *only* by the (directed) representation relation. An intuitive example is given in figure 1: a physical switch is *represented* by an abstract bit, which in this case takes the value 0 for switch state up, and 1 for switch state down. Note, however, that AR theory is not a dualist theory in the sense of Descartes. Everything in the theory is physical in some form. The symbols in the *Abstract* domain in figure 1 are instantiated as ink on paper or pixels on the screen as you read this. What makes them *abstract* in AR theory is that this physical form is to some degree arbitrary, and can change, while still corresponding to the same abstract object.

An example of a physical object in the domain of material entities is a *computer*. It has, usually, internal degrees of freedom, and a physical time evolution that transforms initial input to final output states. An example of an abstract object is a *computation*, which is a set of objects and relations as described in one of the logical formalisms of theoretical computer science. Likewise, an object such as a bacterium is a physical entity, and its theoretical representation within biology is an object in the domain of abstract entities.




Fig. 1. Basic representation has three components: (i) the space of physical objects (here, a switch with two settings); (ii) the space of abstract objects (here, a binary digit); (iii) the directed representation relation \mathcal{R} mediating between the spaces.

The central role of representation leads to the requirement for a *representational entity* (RE). The RE supports the representation relation between physical and abstract. AR theory does not require the RE to be human, or conscious; see [5] for an example of a bacterial RE, which is expanded on in §4 here.

The elementary *representation relation* is the directed map from physical objects to abstract objects, $\mathcal{R}_{\mathcal{T}} : \mathbf{P} \rightarrow M$, where \mathbf{P} is the set of physical objects, and M is the set of abstract objects. We subscript the relation \mathcal{R} with a theory \mathcal{T} to indicate that the relation is theory-dependent. When two objects are connected by $\mathcal{R}_{\mathcal{T}}$ we write them as $\mathcal{R}_{\mathcal{T}} : \mathbf{p} \rightarrow m_{\mathbf{p}}$. The abstract object $m_{\mathbf{p}}$ is then said to be the *abstract representation* (under the given theory) of the physical object \mathbf{p} , and together they form one of the basic composites of AR theory, the *representational triple* $\langle \mathbf{p}, \mathcal{R}_{\mathcal{T}}, m_{\mathbf{p}} \rangle$. The basic representational triple is shown in figure 2(a).

Abstract evolution takes abstract objects to abstract objects, which we write as $C_{\mathcal{T}} : M \rightarrow M$. Again, we subscript with theory \mathcal{T} to indicate that C is theory-dependent. An individual example is shown in figure 2(b), for the mapping $C_{\mathcal{T}}(m_{\mathbf{p}})$ taking $m_{\mathbf{p}} \rightarrow m'_{\mathbf{p}}$. The corresponding physical evolution map is given by $\mathbf{H} : \mathbf{P} \rightarrow \mathbf{P}$. For individual elements in figure 2(c) this is $\mathbf{H}(\mathbf{p})$ which takes $\mathbf{p} \rightarrow \mathbf{p}'$.

2.3 ε -Commuting diagrams

In order to link the final abstract and physical objects, we apply the representation relation to the outcome state of the physical evolution, to give its abstract representation $m_{\mathbf{p}'}$, figure 2(d). We now have two abstract objects: $m'_{\mathbf{p}}$, the result of the abstract evolution, and $m_{\mathbf{p}'}$, the representation of the result of the physical evolution. For some (problem-dependent) error quantity ε and norm $|\cdot|$, if $|m_{\mathbf{p}'} - m'_{\mathbf{p}}| \leq \varepsilon$ (or, more briefly, $m_{\mathbf{p}'} =_{\varepsilon} m'_{\mathbf{p}}$), then we say that the diagram 2(d) ε -commutes.

Commuting diagrams are fundamental to the use of AR theory. If the relevant abstract and physical objects form an ε -commuting diagram under representation, then $m_{\mathbf{p}}$ is a *faithful abstract representation* (up to ε) of physical system \mathbf{p} for the evolutions $C_{\mathcal{T}}(m_{\mathbf{p}})$ and $\mathbf{H}(\mathbf{p})$.

The existence of such ε -commuting diagrams define what is meant by a faithful abstract representation of a physical system. The final state of a physical

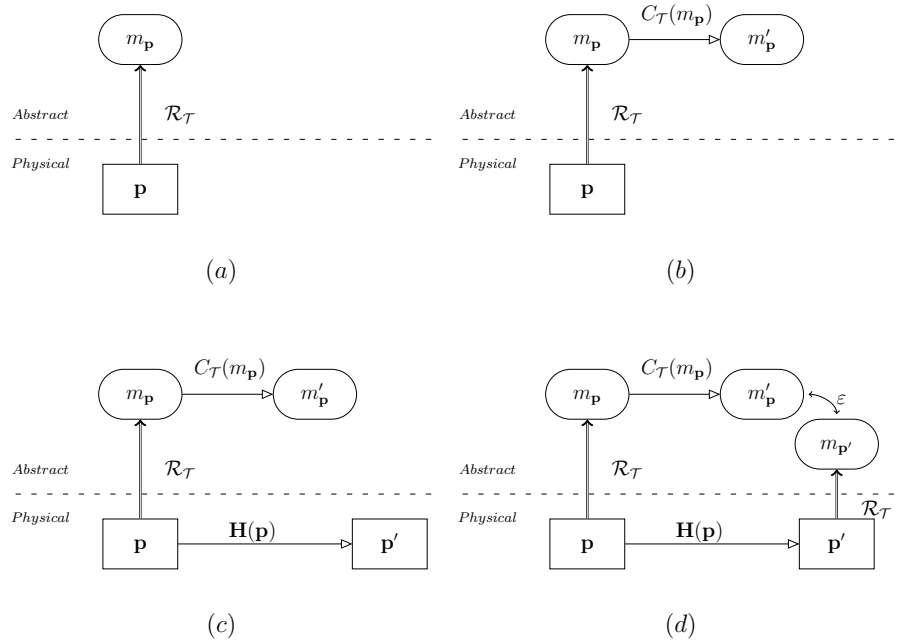


Fig. 2. Parallel evolution of an abstract object and the physical system it represents. (a) The basic representational triple, $(\mathbf{p}, \mathcal{R}, m_{\mathbf{p}})$: physical system \mathbf{p} is represented abstractly by $m_{\mathbf{p}}$ using the modelling representation relation $\mathcal{R}_{\mathcal{T}}$ of theory \mathcal{T} . (b) Abstract dynamics $C_{\mathcal{T}}(m_{\mathbf{p}})$ give the evolved abstract state $m'_{\mathbf{p}}$. (c) Physical dynamics $\mathbf{H}(\mathbf{p})$ give the final physical state \mathbf{p}' . (d) $\mathcal{R}_{\mathcal{T}}$ is used again to represent \mathbf{p}' as the abstract output $m_{\mathbf{p}'}$. If $m_{\mathbf{p}} =_{\varepsilon} m_{\mathbf{p}'}$, the diagram ε -commutes. (Adapted from [3].)

object undergoing time evolution can be known *either* by tracking the physical evolution and then representing the output abstractly, *or* by evolving the abstract representation of the system; and the two results differ by less than the problem-dependent ε . In the first case, the ‘lower path’ of the diagram is followed; in the latter, the ‘upper path’.

Finding out which diagrams ε -commute is the business of basic experimental science; once commuting diagrams have been established they can be exploited through engineering and technology.

2.4 Compute cycle

Figure 2(d) shows the basic ‘science cycle’, of representing a physical system, and determining whether $C_{\mathcal{T}}$ is a sufficiently good abstract model of its behaviour, by requiring that $m_{\mathbf{p}} =_{\varepsilon} m_{\mathbf{p}'}$ for sufficiently many different initial states \mathbf{p} to have confidence in $C_{\mathcal{T}}$ and $\mathcal{R}_{\mathcal{T}}$. There are derived variants of this diagram that capture the ‘engineering cycle’, and the related ‘compute’ cycle. See the original references for details; here we focus on the compute cycle.

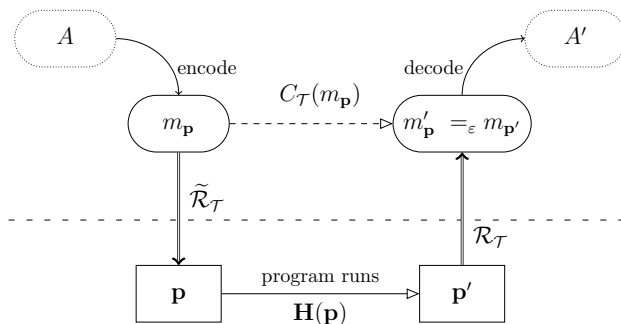


Fig. 3. Physical computing in AR theory. An abstract problem A is encoded into the model $m_{\mathbf{p}}$; the model is instantiated into the physical computer state \mathbf{p} ; the computer calculates via $\mathbf{H}(\mathbf{p})$, evolving into physical state \mathbf{p}' ; the final state is represented as the final abstract model $m_{\mathbf{p}'} =_{\varepsilon} m'_{\mathbf{p}}$; this is decoded as the solution to the problem, A' . The instantiation, physical evolution, and representation together implement the desired abstract computation $C_{\mathcal{T}}(m_{\mathbf{p}})$. (From now on we omit the dashed line separating the physical and abstract world, and rely on the different shaped boxes to indicate what components lie in which domain.)

An ε -commuting diagram in the context of *computation* also connects the physical computing device, \mathbf{p} , and its abstract representation $m_{\mathbf{p}}$. But to do so it makes use of the *instantiation* relation $\tilde{\mathcal{R}}_{\mathcal{T}} : M \rightarrow \mathbf{P}$. Here, instead of saying abstract object $m_{\mathbf{p}}$ represents physical system \mathbf{p} , we say that physical system \mathbf{p} instantiates abstract relation $m_{\mathbf{p}}$. Whereas the representation relation is primitive, the instantiation relation is a derived relation, based on multiple science cycles, abbreviated as $\tilde{\mathcal{R}}_{\mathcal{T}}$; see original references for full details.

The use of $\tilde{\mathcal{R}}_{\mathcal{T}}$ acknowledges that a computer is physical system engineered (or possibly evolved) to have a particular behaviour, rather than a natural physical system being scientifically modelled. The full compute cycle is shown in figure 3, starting from initial abstract problem, through instantiation into a physical computer, physical evolution of the device, followed by representation of the final physical state as the abstract answer to the problem.

Ensuring that the diagram ε -commutes is a process of debugging the physical system, including how it is instantiated (engineered, programmed and provided with input data), and how its output is represented. This shows another key difference from the science cycle: there the diagram is made to ε -commute by instead debugging the abstract model.

The most important use of a computing system is when the abstract outcome $m'_{\mathbf{p}}$ is unknown: when computers are used to solve problems. Consider as an example the use of a computer to perform the abstract arithmetical calculation $2 + 3$. If the outcome were unknown, and the computing device were being used to compute it, the final abstract state, $m'_{\mathbf{p}} = 5$, would not be calculated abstractly. Instead, confidence in the technological capabilities of the computer and the correctness of the instantiation would enable the user to reach the final,

abstract, output state $m_{\mathbf{p}'} =_0 m'_{\mathbf{p}}$ using the physical evolution of the computing device alone. (One advantage of digital computers is that we can achieve $\varepsilon = 0$.)

This use of a physical computer is the compute cycle, figure 3: *the use of a physical system* (the computer) *to predict the outcome of an abstract evolution* (the computation).

2.5 Generality of AR theory

Nothing in the above definition requires the physical computer to be digital, or electronic, or universal, or pre-existing. The computer could be a continuous analogue device; it could be a mechanical or organic device; it could be a hard-wired device with limited capabilities; it could be a ‘one-shot’ device constructed for a particular computation. It simply needs to be sufficiently powerful, sufficiently accurate, and instantiatable, to perform the RE’s desired computations: the relevant squares must exist, and must be known to ε -commute for the desired computations.

And, of most relevance here, nothing in the above definition requires the RE to be a human, or conscious, user. We now show how to model the RE in the same context as the computing system.

3 Including the RE in the model

3.1 Overview

As mentioned above, the representational entity (RE) supports the representation relation \mathcal{R} . Although it does not appear explicitly in the compute cycle of figure 3, it is the *physical entity* that ‘owns’ the abstract problem A and desires the abstract solution A' .

To help clarify the issues, consider a (human) RE who has the problem “I have two apples in my left hand, and three in my right hand; how many apples do I have in total?” We model this physical RE’s problem, how they encode it as a computational problem, how this is instantiated in a physical computer, how the computer finds the answer, how the answer is represented back as an abstract computational result, and how that result is decoded as an answer to the RE’s problem.

In this section, we add the RE to the overall model of physical computing as defined in AR theory. As before, we have objects in two domains: the physical RE, and (our) abstract model of the RE. First we show how we model the RE in a manner analogous to how we model a physical computation (§3.2). Then we show how to integrate the RE and full physical compute cycle models, and how to interpret various parts of the resulting model (§3.3).

3.2 The physical RE

The RE is a physical system \mathbf{p}_{RE} . The relevant part of the RE here is the physical states that it uses to represent its abstract problem A . (There are several levels




Fig. 4. The relationship between the physical representational entity \mathbf{p}_{RE} and the physical computer \mathbf{p}_C via abstract models of each. There is an *encoding* of the abstract model $m_{\mathbf{p}_{RE}}$ into $m_{\mathbf{p}_C}$. In a correctly working system, this encoding is appropriately implemented by the respective physical systems: the square should ε -commute. Note that the models of the RE and the computer are potentially with respect to different theories.

of indirection at play here, that will become clearer with later examples.) Our abstract model of these relevant parts in AR theory diagrams is $m_{\mathbf{p}_{RE}}$; the RE does not in general itself construct AR diagrams. See figure 4. Our model $m_{\mathbf{p}_{RE}}$ may incorrectly capture the RE’s physical state, in which case the model needs to be modified; $m_{\mathbf{p}_{RE}}$ is our *scientific* model of \mathbf{p}_{RE} .

We model the computational system as before. There is the abstract model $m_{\mathbf{p}_c}$ that forms the ‘specification’ of the RE’s problem $m_{\mathbf{p}_{RE}}$ encoded as a computational problem. (This is the model $m_{\mathbf{p}}$ in figure 3.) This is our model of the RE’s model of the computer and encoding: the RE’s model is also part of its physical state \mathbf{p}_{RE} .

The computer’s physical state may incorrectly implement the RE’s model, in which case the physical state needs to be modified; the RE is using an engineering model of \mathbf{p}_c . However, our model $m_{\mathbf{p}_c}$ of the RE’s model may incorrectly represent the RE’s model: we are using a scientific model of the RE and its computer. We model the RE’s problem being encoded into the computer by $m_{\mathbf{p}_{RE}}$ being *encoded* into the computational model $m_{\mathbf{p}_c}$. There is no guarantee that such an encoding is possible: not all problems are computable.

The two representation/instantiation relation arrows in figure 4 are with respect to two different theories. The representation $\mathcal{R}_{T_{RE}} : \mathbf{p}_{RE} \rightarrow m_{\mathbf{p}_{RE}}$ is based on the theory of how the physical RE forms abstract problem specifications; the instantiation $\tilde{\mathcal{R}}_{T_c} : m_{\mathbf{p}_c} \rightarrow \mathbf{p}_c$ is based on the theory of how the physical computer implements abstract computations.

In a correctly implemented computer, the diagram in figure 4 should ε -commute: the instantiated state of the physical computer should correctly mirror the desired state of the physical RE: it should *physically encode* the desired state. The establishment of this physical encoding link is part of the engineering process.

During the execution, this physical encoding link is not necessarily established immediately. There may be some delay, for example in updating a record

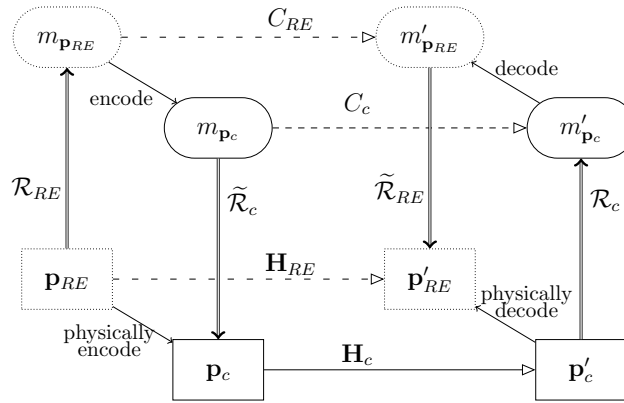


Fig. 5. The full compute cycle including the representational entity and the physical computer. The desired change in the RE’s state, from posed problem to perceived solution, is $\mathbf{p}_{RE} \rightarrow \mathbf{p}'_{RE}$. The physical computer performs $\mathbf{p}_c \rightarrow \mathbf{p}'_c$. The full compute cycle from AR theory is: represent RE’s physical state \mathbf{p}_{RE} (desired computation) as abstract model $m_{\mathbf{p}_{RE}}$; encode to computational model $m_{\mathbf{p}_c}$; instantiate into physical computer state \mathbf{p}_c ; physical computer evolves to final state \mathbf{p}'_c ; represent physical solution as abstract computational solution $m'_{\mathbf{p}_c}$; decode to final abstract problem solution $m'_{\mathbf{p}_{RE}}$; which models the instantiation of the final state of the RE. Each set of squares (between representational entity and physical computer, and across the compute cycle) should ε -commute.

to reflect reality, or in opening or closing a valve to reflect changed demand. In, for example, a mechanical control system, with feedback, there can be an immediate coupling: the behaviour of the physical controlled system changes its state, which is directly communicated to the physical controller through their physical mechanical coupling. We do not consider this aspect further here, although it is a key feature of correctly-engineered computational ‘mirror worlds’ and of feedback control systems.

3.3 The physical RE in the compute cycle

We can now add this physical RE layer to the previous compute cycle. See figure 5 for the full compute cycle including the representational entity. Notice how the RE adds another dimension (cube instead of square) to the diagrams. Each dimension is a level of indirection or representation. The full compute cycle involves traversal of many faces and edges of the displayed cube. Each face has its own place in the model.

Consider again a (human) RE who has the problem “I have two apples in my left hand, and three in my right hand; how many apples do I have in total?”

Back face; RE’s view of the computation (figure 6): the RE’s desired states, starting from a problem state (abstract initial state, $m_{\mathbf{p}_{RE}}$, “how many apples?”; physical state, \mathbf{p}_{RE} , a brain state that is represented by that abstract




Fig. 6. The RE’s view of the problem solution (back face of figure 5). The RE has an initial physical state \mathbf{p}_{RE} , modelled as $m_{\mathbf{p}_{RE}}$. It has a desired final state \mathbf{p}'_{RE} , modelled as $m'_{\mathbf{p}_{RE}}$. Both the horizontal arrows are dashed, as they are implemented in a different medium: the computer.

question) and resulting in a solution state (abstract final state, $m'_{\mathbf{p}_{RE}}$, “five apples!”; physical state, \mathbf{p}'_{RE} , a brain state that captures that abstract solution). There is no direct path from initial to final state, either abstractly or physically, as a separate computer is used to achieve the desired state changes.

Left face; encoding the problem (figure 4): the RE’s initial physical and abstract state encoded into the computer’s initial physical and abstract states. The RE’s abstract problem $m_{\mathbf{p}_{RE}}$ of “how many apples?” can be encoded as the computer’s initial abstract state $m_{\mathbf{p}_c}$ “2+3”. This is instantiated as the computer’s initial physical state \mathbf{p}_c , $\boxed{2+3}$. The RE \mathbf{p}_{RE} physically encodes the problem in the computer’s initial state \mathbf{p}_c by, for example, pressing the keys labelled $\boxed{2}$ then $\boxed{+}$ then $\boxed{3}$. (How this human RE manages to press the keys, given the apples it is currently holding, is an exercise left to the reader.)

Front face; compute cycle (figure 3, which also includes the back face RE abstract models as its ‘abstract problem’ components): the original simple AR theory compute cycle, ignoring the role of the RE. The abstract computational problem $m_{\mathbf{p}_c}$ is instantiated in the computer’s initial physical state \mathbf{p}_c , $\boxed{2+3}$. The physical computer evolves as given by its physical structure, \mathbf{H}_c , which results in the final physical state \mathbf{p}'_c of $\boxed{5}$. This is represented as the final abstract state $m'_{\mathbf{p}_c}$ of “5”. These three steps (instantiation, physical evolution, representation) implement the desired abstract computation C_c : “2+3 = 5”.

Right face; decoding the solution (figure 7): the RE’s final physical and abstract state decoded from the computer’s final physical and abstract state. The computer’s final physical state \mathbf{p}'_c (some kind of pattern of lights in the shape of a figure $\boxed{5}$) is represented as the final abstract state $m'_{\mathbf{p}_c}$ of “5”. This is decoded to the RE’s final abstract state $m'_{\mathbf{p}_{RE}}$ of “five apples!”. The RE’s final physical brain state \mathbf{p}'_{RE} is an instantiation of this, physically achieved by the RE looking at and physically decoding the output from the computer.




Fig. 7. Decoding the solution from the computer to the RE (right face of figure 5). The final state of the computer, \mathbf{p}'_c , is represented as the final abstract state $m'_{\mathbf{p}_c}$; this is decoded to the final abstract state of the RE, $m'_{\mathbf{p}_{RE}}$; and instantiated as the RE's final physical state. This is the model of the physical decoding lower arrow, achieved by the RE physically interrogating the computer.




Fig. 8. The abstract model of the RE's use of the computer to solve its problem (top face of figure 5). The RE has an initial abstract state $m_{\mathbf{p}_{RE}}$; this is encoded into the initial abstract state of the computer $m_{\mathbf{p}_c}$. The computer performs its calculations to produce its final state $m'_{\mathbf{p}_c}$, which is decoded to produce the desired final state of the RE, $m'_{\mathbf{p}_{RE}}$. Both the horizontal arrows are dashed, as they are implemented in a different medium: the physical computer.

Top face; abstract use of a computer (figure 8): the purely abstract view of the (modelled) RE encoding its problem into a (modelled) computation, and decoding the desired solution. There is no direct path from initial to final abstract states as the physical computer is used to achieve the desired abstract state changes. In terms of classical refinement theory [1, 2], C_{RE} can be thought of as the ‘global-to-global’ requirement (although here this need not be captured in a formal manner), with “encode, computation C_c , decode” corresponding to the “initialisation, operation, finalisation” steps.

Bottom face; physical use of a computer (figure 9): the purely physical view of the RE encoding its problem in a physical computer, and decoding the desired solution. That this is a *computation*, rather than some other activity, is established by the abstract models and the various ε -commuting squares.




Fig. 9. The physical system of the RE's use of the computer to solve its problem (bottom face of figure 5). The physical RE has an initial physical state \mathbf{p}_{RE} ; this is physically encoded into the initial physical state of the computer \mathbf{p}_c . The computer evolves over time to produce its final state \mathbf{p}'_c , which is decoded to produce the desired final state of the physical RE, \mathbf{p}'_{RE} .




Fig. 10. The full compute cycle for the bacterial system. See text for details.

All of these relationships must be correctly implemented and modelled (the relevant squares containing encoding, decoding, instantiation, and representation must ε -commute) for the actual physical RE final state \mathbf{p}'_{RE} to be the desired physical RE final state, that is, for the physical computer to have been used correctly, and for it to have performed correctly, to solve the RE's problem.

4 Example: Intrinsic computing in bacteria

Figure 10 shows the RE and the compute cycle in the case of the problem of bacterial computing. This example was originally studied in [5] to demonstrate that it is possible to have computation with a non-human RE. However, without the explicit modelling of the bacterial RE, it resulted in a somewhat circuitous description. With the RE here explicitly present, the model is much clearer.

The physical RE, \mathbf{p}_{RE} , is a bacterium, with a receptor at the front, and a flagellar motor at the back. In the absence of input at the receptor, the motor is off; input causes the motor to switch on, propelling the bacterium towards food. (As ever, the biology is more complicated than this. The original reference should be consulted for the biological details.) The abstract problem, C_{RE} , that

the RE wants to solve is “if there is food, move towards it”. This is encoded as the abstract computational problem C_c : “if there is a signal, switch the motor on”. The abstract signal is instantiated as a particular chemical; the physical *RE* physically encodes the reception of exterior food as the presence of an internal chemical, chem X.

This chemical physically propagates through the bacterium, undergoing transformation via a biochemical pathway, such that another chemical becomes present at the rear. The presence of this other chemical, chem Y, is represented as switching on the (abstract) motor, which is decoded as the answer to the bacterium’s problem: to move. It is physically decoded as activating the flagellar motor.

The bottom face of this bacterial-compute cube shows the purely physical computing: The bacterial RE physically encodes the detection of food by its receptor as a chemical chem X; the biochemical pathway moves and transforms this chemical signal to the rear where it appears as chem Y. The resulting chemical is physically decoded: it attaches to and activates the flagellar motor. The physical problem, of detecting food and moving towards it, has been solved.

That this is indeed a *computation*, rather than a purely physical process, is argued in [5]: chem X, chem Y, and the pathway are in some sense ‘arbitrary’ (they comprise different molecules in different bacterial species), and so it is not their specific *physical* properties, but their *representational*, informational properties, that are being exploited. We are able to model the part of the bacterium that represents the problem as $m_{\mathbf{p}_{RE}}$, and the part that encodes into the computer $m_{\mathbf{p}_c}$ in a way that convincingly contains the right sorts of representation. With representation (and hence a representational entity) identified, we can conclude that there is *data* being processed, not mere *material* being exploited. With ε -commuting diagrams present, we can conclude that computation is present.

5 Conclusion

We have shown how the RE in AR theory can be incorporated into the compute cycle, and how this can illuminate how the physical RE can use a physical device as a physical computer. The RE does not need to be a human brain: the example here shows intrinsic computing by a bacterial RE. This demonstrates how computing, whether conventional or unconventional, can be broader than human use of computers (external or brain-based), but is narrower than pan-computationalism, in requiring the existence of an RE in addition to the computer itself.

It may be that the RE does not even need to be organic, or ‘alive’; it might potentially appear in the loop as an engineered ‘proxy’ for the ultimate RE, for example, the plant in a control system using the controller as a computer to maintain itself in a particular behaviour. In future work we will investigate how far the concept of the RE can be removed from a living user.

Acknowledgements

We thank our colleagues Dom Horsman, Tim Clarke, and Peter Young, for illuminating discussions. VK is funded by UK Engineering and Physical Sciences Research Council (EPSRC) grant EP/L022303/1.

References

1. Clark, J.A., Stepney, S., Chivers, H.: Breaking the model: finalisation and a taxonomy of security attacks. In: REFINE 2005 Workshop, Guildford, UK, April 2005. ENTCS, vol. 137(2), pp. 225–242. Elsevier (2005)
2. He, J., Hoare, C.A.R., Sanders, J.W.: Data refinement refined (resumé). In: Proc. European Symposium on Programming, ESOP 86. LNCS, vol. 213, pp. 187–196. Springer (1986)
3. Horsman, C., Stepney, S., Wagner, R.C., Kendon, V.: When does a physical system compute? *Proceedings of the Royal Society A* **470**(2169), 20140182 (2014)
4. Horsman, D., Kendon, V., Stepney, S.: Abstraction/Representation Theory and the Natural Science of Computation. In: Cuffaro, M.E., Fletcher, S.C. (eds.) *Physical Perspectives on Computation, Computational Perspectives on Physics*, pp. 127–149. Cambridge University Press (2018)
5. Horsman, D., Kendon, V., Stepney, S., Young, P.: Abstraction and representation in living organisms: when does a biological system compute? In: Dodig-Crnkovic, G., Giovagnoli, R. (eds.) *Representation and reality: humans, animals, and machines*, vol. 28, pp. 91–116. Springer (2017)
6. Horsman, D., Stepney, S., Kendon, V.: The Natural Science of Computation. *Comms. ACM* **60**(8), 31–34 (2017)
7. Horsman, D.C.: Abstraction/representation theory for heterotic physical computing. *Philosophical Transactions of the Royal Society A* **373**, 20140224 (2015)