



LEEDS
BECKETT
UNIVERSITY

Citation:

Wilson, AD (2022) Ecological Mechanistic Research and Modelling. *Ecological Psychology*. pp. 1-23. ISSN 1040-7413 DOI: <https://doi.org/10.1080/10407413.2022.2050912>

Link to Leeds Beckett Repository record:

<https://eprints.leedsbeckett.ac.uk/id/eprint/8527/>

Document Version:

Article (Published Version)


Creative Commons: Attribution 4.0

The aim of the Leeds Beckett Repository is to provide open access to our research, as required by funder policies and permitted by publishers and copyright law.

The Leeds Beckett repository holds a wide range of publications, each of which has been checked for copyright and the relevant embargo period has been applied by the Research Services team.

We operate on a standard take-down policy. If you are the author or publisher of an output and you would like it removed from the repository, please [contact us](#) and we will investigate on a case-by-case basis.

Each thesis in the repository has been cleared where necessary by the author for third party copyright. If you would like a thesis to be removed from the repository or believe there is an issue with copyright, please contact us on openaccess@leedsbeckett.ac.uk and we will investigate on a case-by-case basis.



Ecological Mechanistic Research and Modelling

Andrew D. Wilson

To cite this article: Andrew D. Wilson (2022): Ecological Mechanistic Research and Modelling, Ecological Psychology, DOI: [10.1080/10407413.2022.2050912](https://doi.org/10.1080/10407413.2022.2050912)

To link to this article: <https://doi.org/10.1080/10407413.2022.2050912>



© 2022 The Author(s). Published with license by Taylor and Francis Group, LLC



Published online: 16 Mar 2022.



Submit your article to this journal [↗](#)



Article views: 389



View related articles [↗](#)



View Crossmark data [↗](#)

Ecological Mechanistic Research and Modelling

Andrew D. Wilson 

Psychology, Leeds School of Social Sciences, Leeds Beckett University

ABSTRACT

A recent philosophical literature has developed a taxonomy of scientific explanations, models, and the two basic research programmes that produce them. The first programme takes some *capacity* of a system and maps out how it works by breaking it down into various sub-capacities, each with their own distinct characteristics. The end goal is a functional model, a ‘how-possibly’ box-and-arrow type map of the functional organisation of the capacity. The second programme instead focuses on analytically decomposing a proposed mechanism that produces a phenomenon into *real parts and processes*. The end goal is a dynamical mechanistic model, a ‘how-actually’ explanation in which each model part explicitly represents the dynamics of those real parts or processes. Mechanistic models are better explanations of phenomena. Ecological psychology has, so far, widely resisted becoming a mechanistic science. This is in part due to our objections to mechanistic, Cartesian ontologies, and more recently because it’s not clear we can meaningfully decompose the systems we study in order to develop such models. I will argue here that both of these concerns are unfounded, that ecological psychology is actually perfectly capable of developing mechanistic models, and that therefore we should do so, in order to gain the benefits.

Introduction

Scientific programmes aim to develop explanations for various phenomena; why did that thing behave in that particular way when I poked it just now? But not all explanations are equal; some are more explanatory than others. There is now a fairly extensive philosophical literature that proposes a taxonomy of the explanations that get used across science (some exemplar books: Bechtel & Richardson, 2010; Craver, 2007; Craver & Darden, 2013). This *mechanism* (or sometimes *neo-mechanism*) literature lays out, in useful detail, the kinds of explanations science can achieve, from functional ‘how-possibly’ type models that map out the capacities the system needs in order to behave the way that it does, all the way up to fully mechanistic ‘how-actually’ type models that explicitly model real parts and processes of the system. They also describe the kinds of research programmes that produce the various kinds of explanations and discuss when and why you might accept one programme over another. Importantly, while

CONTACT Andrew D. Wilson  a.d.wilson@leedsbeckett.ac.uk; DrAndrewDWilson@gmail.com  Psychology, Leeds School of Social Sciences, Leeds Beckett University, Leeds LH1 3HE, UK.

© 2022 The Author(s). Published with license by Taylor and Francis Group, LLC
This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

mechanistic models are held up as the gold-standard, the literature fully embraces the possibility that not every topic allows the kind of empirical research programme required to produce these models. Functional models come with limits on what they can tell us, but, done properly and with careful attention to those limits, that research programme can be productive and informative and guide future discovery.

Non-ecological cognitive science is currently a busy and productive functional explanation research programme. For many years, the received wisdom has been that this is the best cognitive science can actually achieve, but that this is okay and progress can be made. Bechtel and Abrahamsen (2010) instead challenged cognitive science to rethink this stance, in light of the philosophical work done over the preceding 15 years, and to consider whether cognitive science could, in fact, be mechanistic. If not, fine. But if it was actually possible, then cognitive science would gain access to the explanatory benefits of mechanistic modelling and that could only be a good thing. To help, Bechtel & Abrahamsen walked us through a detailed mechanistic research programme about circadian rhythms to demonstrate how the experiments you run when aiming for a mechanistic model are different, and how you go about *decomposing* a system into candidate parts, *localising* properties to those parts, and then bringing them back together as an explanation of a system being modelled.

Golonka and Wilson (2019a) rose to this challenge and proposed that within cognitive science, only ecological psychology is capable of supporting a mechanistic research programme. Our argument came in two parts; (1) unlike other approaches, ecological psychology trades in real parts and processes, specifically affordances, effectivities, events, and perceptual information variables, and (2) at least one ecological psychology research programme is mechanistic, specifically Bingham's perception-action model of coordinated rhythmic movement (Bingham, 2001, 2004a, 2004b; Herth et al., 2021; Snapp-Childs et al., 2011). We can do mechanisms in principle, we can do mechanisms in practice, therefore we should explicitly aim to do it all the time, in order to gain access to the explanatory benefits mechanisms bring.

There remain two ecological objections, and the goal of this paper is to tackle them head on for an ecological audience. The first is the general ecological distaste for mechanistic approaches to science (e.g. Turvey, 2019). The second, more substantive concern, is whether the ecological commitment to treating the organism-environment behavioural system as the unit of behaviour means we cannot always avail ourselves of the decomposition and localisation empirical strategies required to map out the real parts and processes for modelling (Chemero & Silberstein, 2008; Silberstein, 2021; Silberstein & Chemero, 2013; Stepp et al., 2011). Effectively, the question is, do our ontological commitments make mechanistic research and modelling impossible? I will argue no, and quite the contrary – our ontological commitments actually make us, unlike non-ecological alternatives, perfectly suited to mechanistic modelling. I will suggest that this is great news, and we should embrace the scientific possibilities afforded us by this match.

Functional vs mechanistic research and modelling

The first thing I need to do is lay out the basic rules of functional vs mechanistic research and modelling, and to introduce some key terminology we will need as the

discussion progresses. Because this is a philosophical literature, there are many ongoing arguments about the details, and there are interwoven implications about ‘grand unifying theories’ and the like (Miłkowski et al., 2019). But from the perspective of actual scientific practice, the rules are fairly clear, and the thing I want to emphasise for ecological psychologists is primarily the benefits of the hunt for *real parts and processes*. I’ll cover both functional and mechanistic rules because the compare-and-contrast is helpful for seeing the benefits of the latter, and because I will need to avail myself of this taxonomy in the later sections of the paper.

Functional research and modelling

Functional research and modelling takes as its explanatory target a *capacity* of a system (Craver, 2007). A capacity is simply ‘something that system can do’. The research programme into this capacity then proceeds by decomposing the overall capacity into sub-capacities; other things that the system can do that are needed for the exhibition of the overall capacity. This decomposition continues as long as the data support it, and the end result is a map of capacities of the system that constitute the overall target capacity; a ‘how-possibly’ model.

This will make more sense with a concrete example; I will use the non-ecological cognitive analysis of memory¹, although all the big topics in this psychology (attention, categorisation, etc) work this way. See [Figure 1](#) as we go.

The functional analysis of memory in non-ecological psychology begins by simply noting that the cognitive system seems to have the capacity for operating with respect to things that are not present now but were present earlier. For example, I can tell you what I had for lunch yesterday. This capacity gets the name ‘memory’, and the initial model of this capacity is a single large box labelled ‘Memory’. Now the aim is to map out more of the internal structure of this capacity; what are the sub-capacities that support the overall capacity?

The next thing psychologists showed empirically about the capacity ‘memory’ is that it seems to work on two timescales, and that these have different properties. People can remember things from the very recent past, say up to the last minute or two, but only a limited number of things (around 7, +/-2; Miller, 1956). But they can also remember things from further back, and there doesn’t seem to be any obvious upper limit to the number of things we can remember on this time scale. So now the capacity ‘memory’ is constituted by two sub-capacities; a short-term memory system with hard limits on its capacity (Atkinson & Shiffrin, 1968) and a long-term memory system with no such limits (Tulving, 1972).

So now psychologists experimentally investigate these two systems, and it’s quickly revealed that neither seem to be just one single system. ‘Short-term memory’ gets reframed as ‘working memory’, because it seems to be holding information people are currently using, rather than simply what has happened recently (Baddeley & Hitch,

¹I am using this example solely because it is an excellent and detailed example of a functional research and modelling programme. I can use it to show how such a programme can be productive, informative, and responsive to the evidence, but also how it has hard limit on how explanatory it can be. I do not endorse this analysis of memory, nor am I aiming for a comprehensive review; it’s just an ideal example.

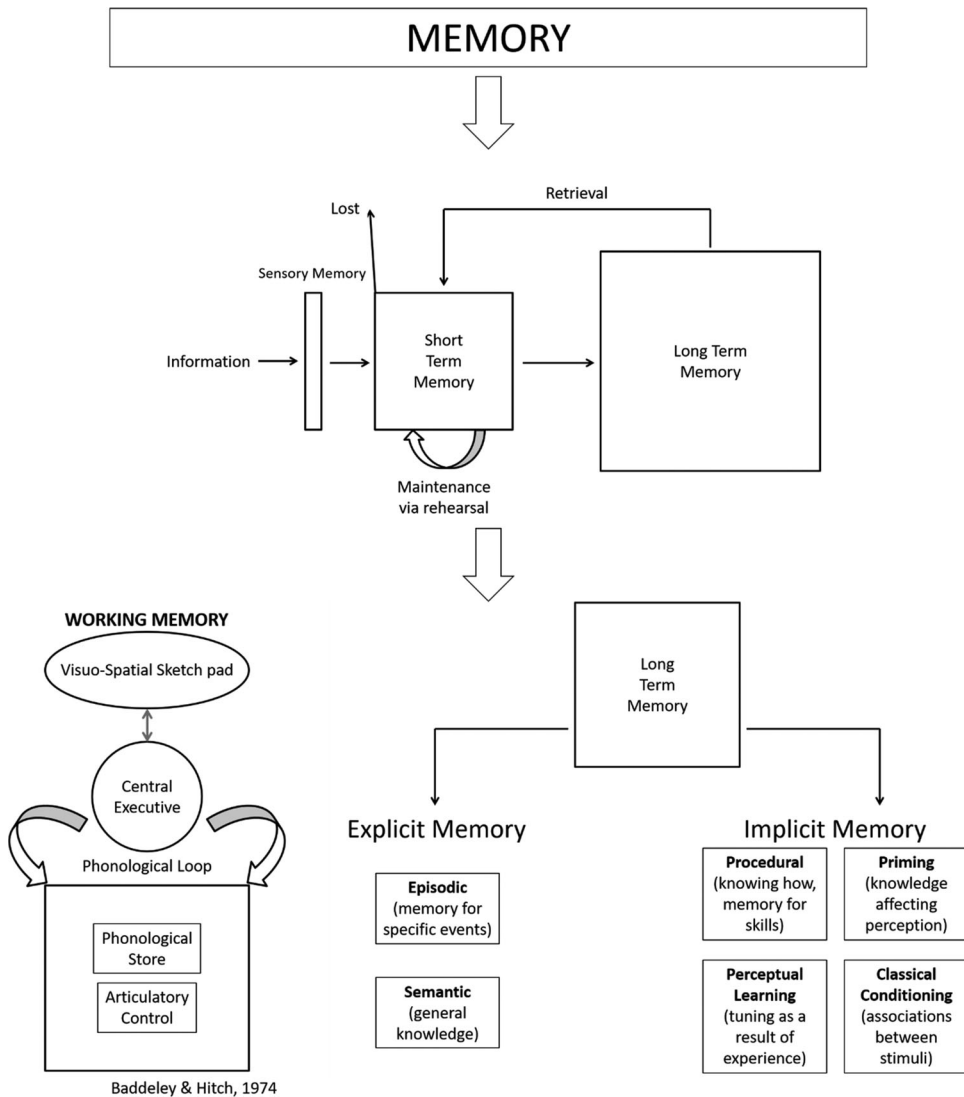


Figure 1. The decomposition of the cognitive capacity ‘memory’ into sub-capacities that seem to constitute the overall capacity.

1974). Within this, there seem to be two main sub-capacities; the ability to work with *phonologically* encoded information and with *visually* encoded information. This is shown by dual-task experiments that reveal when people can and can’t do two things at the same time (e.g. Hitch & Baddeley, 1976). These now need to be coordinated, so a third ‘central executive’ capacity is proposed, and the scientific programme has now decomposed ‘short term/working memory’ into a new layer of sub-capacities. A fourth capacity, the ‘episodic buffer’, has been proposed based on further work, although as yet the evidence remains limited (Baddeley, 2000). ‘Long term memory’ is also constituted by multiple distinct systems, each with their own capacities (Tulving, 1972). People seem to be able to verbalise memories of specific moments in time in one way, a capacity researchers labelled ‘episodic memory’. People also verbalise a large amount of

information that isn't tied to any specific time and place, a capacity researchers label 'semantic memory'. But people can't verbalise everything from their past that can affect our current behaviour, so researchers class those first two capacities as 'explicit memory' and the rest 'implicit memory'. There are at least four kinds of the latter capacity, each with their own distinct capabilities. And, if required, on the research goes.

There are several things to note here, from the point of view of scientific practice. First, this research strategy is highly productive; scientists have been working fruitful careers and know a lot about the functional organisation of memory (e.g. that procedural memory is a type of implicit memory, not explicit). Second, they tell the different capacities apart by the kinds of behaviours they can exhibit, and these are all amenable to modelling (short-term memory modelling in particular is a robust cottage industry).

But third, I also need to note here a limit of the models; they provide a map of the territory, but no hints as to how the territory is actually implemented in a cognitive system. In addition, while these functional maps provide some useful constraints on mechanism proposals, they cannot adjudicate between options that fit the data equally well. Bechtel and Abrahamsen (2010) discuss a great example from the 1980s and 1990s, when symbolic vs. connectionist models were competing as the best architecture to explain the capacities of a variety of cognitive systems. They discuss the extended back-and-forth in the literature about how children acquire the past tenses of regular vs irregular verbs (an example of the capacity 'semantic memory'). Rumelhart and McClelland (1986) proposed a connectionist model that coped with the behavioural data in a single process, and then Pinker and Prince (1988) came back with a dual-process symbolic model that solved the same problem. Each camp then gave the other camp challenges by accounting for some part of the behavioural data the other model didn't currently handle; each time, the camps would adjust their models to handle that case and raise a challenge of their own. This went on for years without any real resolution about which was the 'correct' architecture; the functional models of memory were agnostic about implementation and only required the model to be able to 'save the phenomena' exhibited by that sub-capacity.

Overall then, functional modelling is productive and informative but only to a point, and after that point they simply cannot help constrain our explanations. In particular, because they are agnostic about implementation, they can only roughly constrain any exploration of, say, the neural contribution to a capacity; neural correlates and the like are as far as they go. Non-ecological cognitive science research programmes are, at heart, all functional modelling programmes, and these constraints apply (Charles et al., 2014).

Because they are agnostic about implementation, it is very uncommon and difficult to promote a functional model to a mechanistic one. It is not necessarily impossible, however. To the extent that the various sub-capacities are successfully carving the system at its joints, the functional map might still provide a series of sub-capacities that could each be a target for a mechanistic research programme. The functional modelling might just be a good first swing at developing incomplete mechanism models (mechanistic sketches). The full promotion will still require a different empirical approach, though, which leads us nicely to discussing the rules for mechanistic research and modelling.

Mechanistic research and modelling

Rather than targeting a fairly broad capacity of a system, a programme of mechanistic research begins by identifying a more specific behaviour a system can exhibit², and taking seriously the idea that it arises from the operation of a mechanism, which is here defined as

A mechanism is a structure performing a function in virtue of its component parts, component operations, and their organization. The orchestrated functioning of the mechanism, manifested in patterns of change over time in properties of its parts and operations, is responsible for one or more phenomena.

Bechtel & Abrahamsen, 2010, pg 323

The goal is to model this mechanism, usually using the tools of dynamical systems theory. In order for that model to count as a mechanistic model, however, there is a constraint; each term in the model must explicitly represent the dynamics of a part or process that actually features in the mechanism being modelled. This is the *real components constraint*. The composition and organisation of such a model recreates the composition and organisation of the system being modelled, and therefore stands as a formal expression of the scientific explanation for the phenomena the system produces; a 'how-actually' model.

Craver (2007) notes that while there is no hard empirical threshold for establishing that a part is real, there are four criteria that can help distinguish between real and fictive parts. First, real parts have a stable cluster of properties; this suggests that the part is a non-accidental feature of the system. Second, real parts are robust; they can be identified with multiple methods, which suggests that the part isn't an accidental feature of the measurement method. Third, real parts can be used to intervene with other components; they are causally related to the operation of other components. Fourth, real parts are plausible given the system being studied; physiologically, or physically, or otherwise. These are all heuristics, but the general idea is a part can count as real once you have accumulated consistent evidence for what it does and how it affects other parts using multiple methods across numerous contexts.

The real components constraint therefore means that the scientist needs empirical evidence that the various components exist, what dynamical properties they exhibit, and how they play a part in the mechanism being modelled that relates to those dynamics. Gathering this evidence initially requires a process of *decomposition* and *localisation*. Decomposition is the analytic process of identifying candidate parts, and empirically testing those parts to identify their dynamical properties (effectively establishing whether those candidates are, in fact, real parts). Localisation is the analytic process of attributing those properties to actual parts of the mechanism.

Decomposition is generally a heuristic process; you try breaking the mechanism into parts according to some scheme and those parts either work in ways that contribute to the mechanism, or you have the wrong parts and must try again. For example, I can decompose an alarm clock with a hammer or a screwdriver in order to find out why it rings loudly at the same time every morning; the former is fun, but it produces pieces

²As an aside, I also think the task specific focus of mechanistic research and modelling is a more natural fit for ecological psychology as compared to the general capacity focus of the functional approach.

that, upon testing, don't feature in the production of the phenomenon to be explained. Finding the scheme that produces the right parts is actually very hard³, but it identifies the crucial *grounding scale* for the mechanistic model (Bechtel & Abrahamsen, 2010; Golonka & Wilson, 2019a). This is the scale at which the phenomenon lives, and it is the scale that all the activity of the parts and processes are organised with respect to. For example, the phenomenon of fermentation lives at the biochemical scale, and is best explained by a model using biochemical molecules as real parts (this example from Bechtel & Abrahamsen, 2010). Of course, biochemical molecules are themselves constituted by smaller molecules, and in turn by various atoms; in addition, fermentation happens in a larger context. Efforts to use either the lower or higher scale parts to model the mechanism of fermentation per se failed, however, and this failure revealed these were the wrong grounding scales; they weren't providing the right parts to explain fermentation. These other scales *can* provide parts used to build parts that feature in fermentation in a multi-scale mechanistic model, but what we use those other parts for is dictated by the needs at the correct biochemical grounding scale⁴.

Decomposition and localisation must therefore take place with constant, explicit reference to the proposed grounding scale and the function of the mechanism. Trying to ground a mechanistic model at the wrong scale produces parts that don't feature in the function of the mechanism being investigated, and this is considered evidence for changing the decomposition and localisation work. The correct grounding scale is rarely, if ever, the smallest scale; mechanistic modelling is very explicitly not a reductionist approach. Quite the contrary; while modelling the mechanism at the grounding scale may be all that's required for a satisfactory explanation, additional scales can also be modelled, and the relationship between the scales modelled as well. In addition, part of the empirical process has to be resituating the mechanism back in its typical context (Bechtel, 2009). Good mechanistic models are therefore typically multi-level, or multi-scale⁵.

Mechanistic research produces better models

This taxonomy of scientific research programmes and the type of explanations they lead to has a normative element: the claim is that the models that come from a mechanistic research programme are scientifically better than ones that come from functional research programmes, and so, if available, this research and these models should be preferred. Bechtel and Abrahamsen (2010) frame this by listing six scientific tools mechanistic models give you. They: (a) demonstrate that a given mechanism is sufficient to produce the target phenomenon, (b) explore the functioning of the mechanism in a larger parameter space than is accessible in experiments, (c) identify whether candidate parts are essential to the mechanism's functioning, (d) explore how particular types of

³I'll note here the ecological ontology provides an excellent decomposition guide for cognitive systems, which is a genuine strength we bring to mechanistic modelling in cognitive science.

⁴Also, whether this is work worth doing is essentially up to the scientist; the biochemical scale analysis is perfectly mechanistic by itself and may also be as explanatory as required.

⁵The word 'level' is a tricky one, and while it generally communicates the point a speaker wants to make, it does tend to cause trouble in more formal analyses. This has led several authors to prefer the word 'scale' (e.g. Eronen, 2015) as a more defensible way to express this idea, so I will rely on 'scale' from here on out. I just mention 'levels' here because that gets used in the literature and I wanted to make the link clear.

damage might affect the system by perturbing the model in particular ways, (e) explain how coordinated behaviour can emerge from the coupling of simpler mechanisms, and (f) explore the consequences of altering the relations between multiple mechanisms. Functional models have access to the first four benefits, but only mechanistic models provide access to the full list. That said, while mechanistic models are considered to be best practice, good functional modelling also has its place and the neo-mechanist literature does not demand we only do mechanisms; “While our view is conservative, it is not imperialistic. Specifically, we do not intend 3M [model-to-mechanism mapping] to rule out nonmechanistic explanation generally” (Kaplan & Bechtel, 2011). If it turns out that there are good reasons why ecological psychology cannot be mechanistic, and if there is a good alternative research and modelling strategy available, this will be ok.

To summarise; mechanistic research analytically decomposes the mechanism responsible for the phenomena of interest into candidate real parts and processes, and then empirically characterises the dynamical properties of these parts and processes to see whether those properties are the ones that come together to produce the phenomenon. If not, this is a hint you are decomposing according to the wrong scheme or at the wrong scale. If yes, then you can bring these parts and processes back together into a model, organised with respect to each other in a way that reflects their organisation in the mechanism-being-modelled. You then gain several key scientific advantages over functional research and modelling. For example, the real components constraint means you can do meaningful science *on the model itself*, via simulations; perhaps perturbing or lesioning the model of the system in ways you cannot do to the system being modelled. This then feeds directly into future research aimed at expanding and improving the fidelity of the model representation of the actual system. If you can have them, dynamical causal mechanistic models are pretty great and worth pursuing.

The next question is therefore, can ecological psychology do mechanistic research and modelling? In the next section, I will argue that we can, by addressing the main ecological objections.

Ecological objections to mechanistic research and modelling

I noted at the start that there are two related concerns about whether the ecological approach can or even should embrace mechanistic research and modelling.

The ecological ontology is non-mechanistic

The first concern, that the ecological ontology is explicitly non-mechanistic, is easily addressed. The ecological ontology is actually explicitly *non-Cartesian* (Gibson, 1979; Turvey, 2019); ‘Cartesian’ and ‘mechanistic’ used to be synonyms but aren’t any more (hence ‘neo-mechanists’). Ecological psychology rejects the idea that the only real parts and processes we have to work with are those of 17th century Cartesian physics, specifically primary and secondary properties. We also reject the Cartesian claim that the only notion of causation we have is linear chains of force transmission.

Fortunately, the neo-mechanistic literature is not committed to a Cartesian ontology (Kaplan & Craver, 2011). In fact, it is not especially committed to any particular

ontology, other than the basic realist assumption that there really are parts and processes out in the world to be found. This is because the neo-mechanistic literature is primarily an epistemological project, not an ontological one. It does not specify *what* your candidate real parts and processes should be; it only specifies that your *explanations* should make explicit reference to real parts and processes, and proposes empirical criteria for what counts as real (Craver, 2007). What those candidate real parts and processes are has to be provided by a relevant theory that includes an ontology. Any ontological system that defines a set of real parts and processes that can be empirically shown to be up to the job of producing the phenomenon at hand could, in principle, produce a mechanistic model of that phenomenon. In fact, trying to mechanistically model using various sets of allowed real components is actually a good way to see if your ontology is up to the challenge. Piccinini (2020) is trying to do this with the ontology of computational representations. Ecological psychologists reject this Cartesian ontology because we have reasons to believe it is *not* up to the job of explaining intentional, perceiving-acting physical systems. Instead, we flipped the question and ask ‘what kind of real parts and processes *can* produce such systems?’ and we end up with the ecological ontology of candidate real parts and processes such as dynamical affordances and effectivities, events, and kinematic perceptual information variables that specify these, plus more interesting notions of causation from complexity science and modern physics (Turvey, 2019). Regardless of whether it will work in practice, ecological psychology is at least well-placed to consider mechanistic research and modelling because we have an ontology of candidate parts and processes.

You can't decompose an organism-environment behavioural system

The second concern is about whether it will, in fact, work in practice. The key to developing a mechanistic explanation and model of a phenomenon is the analytic and experimental processes of decomposition and localisation; identifying the real parts and process operating at the relevant scale, characterising their dynamical properties, and figuring out where in the dynamical organisation of the mechanism each piece lives. This heuristic process can fail in two circumstances. First, it fails if there are no parts to be found; connectionist networks are the typical example here. Second, and more relevant to ecological psychology, is that it fails when there are real parts and processes, but these interact nonlinearly in the mechanism that produces the phenomenon. The implication of this nonlinearity is that the dynamical properties of a part in isolation are not, in general, the dynamical properties of the part when it is in the mechanism. The mechanistic research programme of studying isolated components therefore fails to help develop an explanation of the intact system.

The ecological approach is broadly committed to the idea that behaviour emerges from nonlinear organism-environment systems, coupled via information. The practical upshot of this is that we must study perception-action systems as *systems*, and not in terms of isolated parts. The field seem therefore to immediately fall foul of the second obstacle to mechanistic research and modelling (Chemero & Silberstein, 2008; Silberstein, 2021; Silberstein & Chemero, 2013; Stepp et al., 2011).

There are two options available to us, if this is true. First, we can settle for functional analyses of behaviour, which are actually pretty informative, just limited. Second, we could see if there are other ways to *explain* besides mechanistic modelling. Chemero and Silberstein (2008) and Silberstein and Chemero (2013) adopt this second strategy, and propose that *dynamical* research and modelling can produce perfectly good explanations beyond functional analyses. They advocate ecological psychology to adopt this strategy. I have two responses to these points for our field to consider; first, dynamical models do not explain, and second, done carefully, it is, in fact, possible to meaningfully decompose and localise the real parts and processes of a perception-action system.

Dynamical models do not explain

The argument in favour of dynamical explanations is that dynamical models can get explanatory power from their ability to successfully predict phenomena, even if they make no reference to the underlying causal structure (Chemero, 2009; Chemero & Silberstein, 2008; Menary, 2007; Silberstein, 2021; Silberstein & Chemero, 2013; Stepp et al., 2009). There are three basic concerns that have been raised in response (Kaplan & Bechtel, 2011; Kaplan & Craver, 2011).

First, it's actually quite trivial to show that prediction and explanation are not the same thing; for example, I can predict that a storm is forming based on falling mercury in a barometer, but that falling mercury does not explain why the storm is forming. Second, if predictive power and explanatory power are the same, there is no way to adjudicate between equally predictive models that rest on very different assumptions (as with the argument about verb tense learning discussed above) or include irrelevant details. The history of science is full of examples of this mattering; for example, both Ptolemy's geocentric model of the solar system and Copernicus' heliocentric models can *save* the phenomena but it is important that we are able to identify that only the latter *explains* those phenomena. Third, if predictive power and explanatory power are the same, there would be no scientific benefit from developing more mechanistic models; Copernicus would have had no motivation to overturn Ptolemy.

The first point is a concern about whether prediction has the necessary link to explanation; the second and third points are about how scientific progress would be hampered if prediction and explanation are equated. The net result is strong reason to think that dynamical covering law models, while interesting and powerful, do not get any explanatory power from their ability to predict (Kaplan & Bechtel, 2011; Kaplan & Craver, 2011).

I would like to add a concern, specifically that dynamical models can be seen to fail very quickly to predict characteristics of systems beyond the scope of the current dynamical equations. Take coordination dynamics. Chemero in particular has appealed to this topic on multiple occasions when arguing that dynamical models can be suitably explanatory (e.g. Chemero, 2009; Stepp et al., 2009). In particular, of course, he appeals to the history of the Haken-Kelso-Bunz style research and modelling of bimanual coordinated rhythmic movement (Kelso, 1995). This model is literally the exemplar of the application of dynamical systems theory to behavioural systems. It has been extraordinarily influential and pushed the development of dynamical models of behaviour in ways

that ecological psychology has greatly benefitted from, and it deserves its prominent place in the history of our field.

However, it also provides a clear illustration of why dynamical models describe, and do not explain. The HKB model accounts for the particular organisation of bimanual coordinated rhythmic movement via an attractor landscape. Specifically, the equations (organised at the scale of the order parameter relative phase) make it so there is a deep, steep attractor centred on 0° , a less deep, less steep attractor centred on 180° , and a repeller centred at 90° . The equations don't explain *why* bimanual coordinated rhythmic movement has this attractor layout; it simply describes the fact that it does have this attractor layout (Bingham, 2004a). This became obvious when Zanone and Kelso (1994) made several predictions about how learning would proceed, based on the characteristics of the attractor landscape described in their equations of the untrained system; effectively, trying to use their characterisation of the intrinsic dynamics to predict dynamics not currently summarised in the equations. They predicted that relative phases at some close distance to 0° (e.g. 30°) would be harder to learn than relative phases at that same close distance to 180° (e.g. 150°), because the attractor at 0° was deeper and steeper and would therefore more strongly attract any attempt to perform relative phases close by, compared to the less deep, less steep attractor at 180° . This was empirically tested in two experiments (Fontaine et al., 1997; Wenderoth et al., 2002) and these studies found that precisely the opposite was true. Wenderoth et al suggested, and later work by Bingham (discussed below) confirmed, that the reason is the region around 0° is perceived very clearly, making 30° readily discriminated from 0° , while the region around 180° is perceived less clearly, making 150° relatively harder to discriminate. All this work pointed to the fact that *0° is not easy because there is a strong attractor there; it can be described as having a strong attractor there because it is easy.* We still need a perception-based explanation of why it is easy, and the dynamical HKB model cannot serve as a guide to discovery to find that.

So, while dynamical systems modelling is excellent and a powerful part of the ecological research programme, the resulting models are not explanatory because the attractor dynamics they lay out merely summarise the data. As a result, attempts to predict dynamics beyond that data tend to fail. Attractors are not real parts, and it turns out that matters for our ability to 'derive, predict and discover important things' (Silberstein & Chemero, 2013, p. 964). This means that if nonlinear organism-environment behavioural systems really cannot be meaningfully analytically decomposed into real parts we can localise properties to, dynamical descriptions and functional analyses are the best we can get; no explanations for us. This would be okay, but a shame.

Luckily, this isn't the case.

You can analytically decompose an organism-environment behavioural system

To illustrate this point, I will use Bingham's perception-action model of coordinated rhythmic movement that he and colleagues (including myself) developed to replace the merely descriptive HKB model. This is an ecological programme of research that, uniquely, adopted the mechanistic research strategy of decomposition and localisation, and the resulting model stands as a better, 'how-actually' model of the mechanism underpinning the phenomena of this task. I want to use this programme as an example

of what such a programme looks like in practice, and how the discipline of hunting for real parts and processes paid off in terms of explanatory power. I will also draw the reader's attention to how that programme was always guided with explicit reference to the behaviour of the mechanism at the system scale, which underscores the importance of a good grounding scale and shows how to handle the problem of nonlinear interactions and the loopiness of the systems we study. This is how the ecological approach helps guide what Bechtel (2009) calls 'looking down, around, and up' in order to keep mechanistic explanations situated.

The action parts. In bimanual coordinated rhythmic movement, two human limbs oscillate at some frequency and amplitude. The first question is therefore, *what kind of real part is a rhythmically moving human limb?* Note that we are not asking about limbs, but rhythmically moving limbs, because this is how the parts 'limbs' will feature in the full mechanism producing the phenomena of bimanual coordinated rhythmic movement.

Two perturbation studies (Kay et al., 1987, 1991) took single rhythmically moving human limbs, initially moving at a steady state with a set frequency and amplitude. They then mechanically perturbed them to identify their dynamical properties, by scaling the required frequency or by briefly speeding or slowing the motion with an external force. The authors then quantified how the system went about returning to its steady state with the aim of identifying what type of dynamical properties the intact system exhibited. They found five key dynamical properties; limit cycle stability, phase resetting, an inverse frequency/amplitude relation, a direct frequency/velocity relation, and a rapid relaxation time that is independent of frequency. These properties mean that the dynamics of the oscillating human limb are those of a *nonlinear damped mass-spring* driven by an *autonomous* driver (one composed of state variables, not by an explicit function of time). The basic form of a damped mass-spring is

$$\ddot{x} + b\dot{x} + kx = 0$$

where \ddot{x} , \dot{x} , and x are the *state variables* position and its derivatives velocity and acceleration, and b and k are *parameters* (damping and stiffness, respectively). You keep such a spring moving by *driving* it (keeping the equilibrium point equal to something other than 0), and you drive such a spring *autonomously* with a driver composed only of state variables.

Kay et al. (1987) modelled the limb by combining two well-known systems (the Rayleigh and the van der Pol oscillators) into a customised hybrid. This hybrid oscillator was just a dynamical summary of the relevant data, though, and not a model of an actual limb. The best candidate dynamical model of actual limbs is the physiologically motivated λ -model (Feldman, 1986). When Bingham came to model the coordination task, he instead two λ -model style damped-mass springs and drove them with their own phase (based on Bingham, 1995, 2004b; Goldfield et al., 1993; Hatsopoulos & Warren, 1996);

$$\ddot{x} + b\dot{x} + kx = c \sin(\theta)$$

where

$$\theta = \arctan(\dot{x}_n/x), \quad \dot{x}_n = \dot{x}/\sqrt{k} \text{ and } c = c(k)$$

The model copes with the perturbation results by modelling the stiffness k as a function of the sensed departure from the limit cycle

$$k = k_i + \gamma |e_i e_i|, \quad e_n = \sqrt{(v_n^2 + x^2)}$$

This phase-driven oscillator fulfils all the criteria established above (the full oscillator has limit cycle stability and is driven autonomously) and produces all the observed dynamical properties (phase resetting, an inverse frequency/amplitude relation, a direct frequency/velocity relation, and a rapid relaxation time that is independent of frequency). This, I suggest, is a dynamical model of a real action part of the mechanism responsible for the characteristics of bimanual coordinated rhythmic movement, and we've been able to characterise its dynamics outside of the full bimanual coordinated movement task; decomposition and localisation were possible, and they worked.

A couple of things to note here. First, we have not decomposed the system very far; we have so far moved only from the two coupled rhythmically moving limbs to looking at the dynamics of a single rhythmically moving limb. Second, at this stage, we have studied the single limb system with perturbation methods, which keep that system intact but alters candidate state variables to see how the system responds. However, we are still doing decomposition and localisation relative to the full coordinated bimanual system, and we have identified a real part (a phase-driven damped-mass spring) whose dynamical properties have been localised to specific elements. This is a mechanistic approach!

Third, Bingham stopped at this stage because this was as far as he needed to go to identify a relevant real part. As I said above, where you stop is up to the needs of the scientist. If anyone wanted to go further, and identify the neuromuscular parts and processes that can get softly assembled into a rhythmically moving limb, they can, and they will do well so long as they proceed relative to this grounding scale. Such a programme exists, of course, in the form of Feldman's physiological scale equilibrium-point hypothesis, and I'll discuss at the end the possibility of an ecological neuroscience scale addition to all this as well.

The next step was to take two such oscillators, and couple them via an information variable.

The perception parts. A rhythmically moving limb is a real part that is, dynamically, a damped mass spring driven autonomously by its own perceived phase. To implement part of the coupling, each limb when coordinating with another limb is modelled as a damped-mass spring driven autonomously by the perceived phase of the *other* limb. All that was left was to identify the component of the driver that was relative phase specific and that produced the HKB phenomena (0° easy and stable, 180° harder but still stable, 90° maximally unstable).

This component is also perceptual, not neural or neuromuscular (the HKB phenomena persist between people watching each other; Schmidt et al., 1990). The question is therefore, how is the dynamical property *relative phase* specified in information? The information variable needs to be created when two rhythmically moving limbs are present, it must specify relative phase, it must be constituted only by state variables of the damped mass spring (to preserve the overall autonomy of the dynamic) and it must

produce the HKB phenomena, which it turns out only emerge when the oscillators are moving in parallel with each other (Bogaerts et al., 2003; Wimmers et al., 1992).

Bingham and colleagues ran a series of judgement studies in which they studied the perception of relative phase in isolation from the production of bimanual coordinated rhythmic movements (a decomposition and localisation approach). They first presented participants with visual point light displays of coordinated rhythmic movements, using either recorded movements or simulations. They also produced these motions moving side-to-side or in depth (Bingham et al., 1999; Zaal et al., 2000) and they scaled the frequency of the motions (Bingham et al., 2001). Finally they also tested the haptic perception of relative phase, by moving people's limbs (Wilson et al., 2003). Across these studies, they had people judge the mean relative phase or the phase variability.

The results identified more key characteristics the variable specifying relative phase had to have. It was detectable in the same way via vision and haptics. In both cases, it was best discriminated around 0° , less well around 180° , and not at all around 90° . Increasing the frequency reduced discriminability, but mostly at 180° . The variable is defined throughout the trajectory; people could still make these judgements with parts of the motions occluded, although occluding the ends made the task harder than occluding the middle (Bingham, 2004a).

There are several candidate information variables for relative phase that are created by the presence of two rhythmically moving limbs, specifically relative speed, relative frequency, relative position and relative direction. The results above all favoured relative direction as the specifying information; of all of the candidates, it was the only one with all the necessary properties. To confirm this, Wilson and Bingham (2008) perturbed each variable separately in visual displays of rhythmically moving dots and had people judge them in terms of their coordination. The perturbations altered the motion of the displays in ways that kept all the candidate variables present and preserved the relative phase, but systematically stopped each candidate variable from specifying relative phase. If a person was using that perturbed variable to perceive relative phase, judging the relative phase would become impossible. Relative direction proved impossible to perturb without making the displays stop being coordinated rhythmic movements. Perturbations of relative speed, relative frequency and relative position simply added noise to the judgments, and did not stop people from doing the task⁶. Relative direction is in the information variable most people use to perceive relative phase (at least, prior to training). We can now couple our (real) phase driven oscillators via our (real) informational variables.

A couple of things to note. Again, we haven't decomposed the system far; people are judging properties of coordinated rhythmic movement displays. But again, we have studied a decomposed component (relative to the intact bimanual production system), and localised certain properties of that intact system (specifically, the HKB pattern) to the information specifying relative phase. And again, all the perceptual parts are real features of the optic or haptic array.

⁶3 participants were perturbed by the position perturbation at 0° and 180° , and all participants trained to perceive 90° were perturbed by this as well, suggesting they were perceiving relative phase using relative position. This issue of individual variation and the consequences of learning are part of the next, ongoing round of research, discussed below.

The mechanistic model of bimanual coordinated rhythmic movement. The result of the mechanistic research programme is a model of the intact perception-action mechanism producing bimanual coordinated rhythmic movement prior to any training, built of real parts identified in experiments on the parts that characterised their dynamical properties (Bingham, 2001, 2004a, 2004b);

$$\ddot{x}_1 + b\dot{x}_1 + kx_1 = c\sin(\Phi_2)P_{ij}$$

$$\ddot{x}_2 + b\dot{x}_2 + kx_2 = c\sin(\Phi_2)P_{ij}$$

where $P_{ij} = \text{sgn}(\sin(\Phi_1)\sin(\Phi_2)) + \alpha(\Phi_1 - \Phi_2)N_t$

The model is two phase driven oscillators, coupled by the perceived phase of the other. This coupling is modified by the perceived relative direction, which is implemented as the relative direction of motion, discrimination of which is conditioned on the relative speed (experimentally confirmed by Snapp-Childs et al., 2011 in another mechanistic experiment on the decomposed task). This model fits all the criteria and reproduces all the phenomena, across both movement and judgement experiments.

Because the model is made of experimentally verified real parts and processes, the model has also had successes in accounting for phenomena beyond the tasks it was designed to model. Specifically, it has been guiding work on how other relative phases besides 0° and 180° are learned. Because in the model poor performance at 90° is caused by informational coupling, and not by other factors such as muscle homology, we predicted that learning to perceive 90° clearly would lead to improved performance at 90° , without training at performing 90° ; this was confirmed by Wilson et al. (2010). Then, because relative direction is intrinsically variable at 90° , this suggested improved performance there might require learning to perceive relative phase via a different variable; that variable is always relative position (Leach et al., 2021a; Snapp-Childs et al., 2015; Wilson & Bingham, 2008). Relative position also supports perception of relative phase across the board, including 0° and 180° (Wilson & Bingham, 2008) and 60° (Leach et al., 2021), and this supports the transfer of learning 90° or 60° to other relative phases (Leach et al., 2021a, 2021b).

These are all notable in contrast to the failure of the HKB attractor dynamics account to serve as a guide to discovery beyond the data it summarised. Bingham's model above does not model learning or the resulting learned dynamic, but it has still served (along with a focus on only allowing real parts to play a role) to guide experimental work on learning for 15 years. The results of that work have now been implemented in a new version of the model that includes learning, which, as before, contains only real parts and processes (Herth et al., 2021). It now makes some further predictions about the process of learning that will require more experiments, and so on it goes.

Summary of the responses

The ecological objections to mechanistic modelling are that our ontology is explicitly not a mechanistic one, and that we study the type of nonlinear systems that resist decomposition and localisation. As a result, it has been argued that we should focus our attention on developing dynamical models instead, although that this is ok because these explain phenomena pretty well.

This section has provided some replies, specifically 1) our ontology is not the mechanistic Cartesian one of primary and secondary qualities, but rather a mechanistic ecological one of dynamical affordances and effectivities, events, and kinematic information variables, 2) dynamical systems models are great but they do not explain the way mechanistic ones do, and 3) mechanistic research and modelling is an option for ecological psychology because at least some decomposition and localisation is, in fact, possible and effective using our parts. Because mechanistic research and modelling produces top tier explanations, and because ecological psychology can do mechanistic research and modelling, I propose that we embrace this approach and see where we can take it.

The mechanism taxonomy and ecological research

I've argued here and in Golonka and Wilson (2019a) that the ecological approach can be mechanistic because, uniquely in the cognitive sciences, our ontology and research activity already explicitly trades in real parts and processes (affordances, effectivities, events, and information). Perhaps we are already doing the good science and we don't need to worry about this particular philosophical literature. In this section, I want to justify why I think explicitly adopting the neo-mechanistic mindset and empirical practices is worth our attention; the take-home message is that while our ontology includes these parts, our models mostly do not, but they are trying to and the mechanism literature can help guide that development.

Evaluating current ecological modelling within the taxonomy

While ecological psychologists deal in real parts and processes in our theory, modelling work is more variable. Bingham's work is unique in explicitly adopting a mechanistic programme of decomposition and localisation. The field is actually dominated by two basic approaches.

Dynamical covering law models

These are models like the HKB model, which contain no reference to any real parts or processes but instead summarise the dynamical characteristics of the system being studied in terms of attractors and repellors. These are not specific to a particular implementation of the system; the HKB dynamics, for example, have been applied to a wide variety of situations, from bimanual coordination to interpersonal coordination (e.g. Schmidt & Richardson, 2007) but also social memory (e.g. Nordham et al., 2018) and social behaviour in general (e.g. Tognoli et al., 2020), speech (e.g. Nam et al., 2009; Port & Kaipainen, 2002) and brain activity (e.g. Frank et al., 2000). This wide applicability is possible because the model is of higher-order dynamical patterns, and not specific real parts and processes. This is a powerful and informative approach, but as noted above, such modelling cannot explain why the patterns are present, only describe that they are present. Useful, therefore, but limited, and done with no reference to ecological real parts and processes.

Behavioural dynamics models

Probably the most common form of research and modelling in ecological psychology can be broadly referred to as *behavioural dynamics* (Warren, 2006). Across all the many examples, this modelling strategy does analytically decompose the organism-environment behavioural system one step; into agent dynamics, and information about the environment. Information variables are then typically modelled in a real parts fashion; the model includes a specification of the actual information variable hypothesised to be used in performance of the task. However, this is as mechanistic as the model and research strategy gets. Agent dynamics are typically modelled with the kind of abstract dynamical systems models seen above; the goal is simply to describe the observed attractor/repellor landscape. In addition, the empirical research always takes place on the intact system; there is no recourse to any process of decomposition and localisation. These models do gain some explanatory power relative to dynamical covering law models, thanks to the inclusion of terms that represent an actual real ecological part, but only some power, as I'll now illustrate.

Fajen (2007) calls the resulting models examples of *information-based control*. Because they do not model organism-agent dynamics with the real ecological parts 'effectivities' and 'affordances', they run into a problem. They rely on information variables specifying current states of the agent (e.g. "I am (not) currently braking hard enough to stop") rather than about affordances and effectivities (e.g. "it is (not) still possible to stop within the limits of my braking capabilities."). As a result, they do not respect *action boundaries*; the model has no information in it about whether it is possible for the agent to achieve the information-based control goal. They therefore fail to explain many features of skilled action which entail sensitivity to action boundaries. Fajen advocates moving towards *affordance-based control*, in which the information is constrained to be of affordance properties, and models include a consideration of the agent's effectivities; the result is a better explanation of actual behaviour in his tasks. The result is a move towards explicitly representing more ecological real parts and processes in our models; it has, at least, mechanistic motivations. Again, however, the research strategy focuses on studying the intact system only (no decomposition and localisation) and so it is not yet an explicitly mechanistic programme of research and modelling; it has also not become very common yet.

Ecological psychology remains dominated by behavioural dynamical research and modelling, but it is also clearly engaging with questions of whether these count as an adequate explanation, and how we can include the real parts and processes from our ontology in our models. These questions are precisely what the mechanistic literature is designed to provide answers to; research via decomposition and localisation, and modelling of the resulting real parts and processes. If we engage with this literature to attempt to explicitly develop dynamical mechanistic models, we will have useful guidance as we tackle important issues in ecological psychology and work to produce the kind of better models researchers want, and which Bingham's work show us are possible.

Informing future ecological research with the taxonomy

Besides offering a framework to guide the development of research programmes and models that explicitly align with the ecological ontology, I suggest here that the real

components discipline of the mechanistic modelling strategy can help inform theory development and help preserve our rigour as we work to expand the scope of the ecological approach.

Neuroscience

One key domain in which the real components constraint can provide crucial discipline is in developing a Gibsonian neuroscience⁷ that can be a real contender in the field. As a field we have sensibly avoided jumping into studying the brain, as we recognise that we must first provide a behavioural scale ‘job description’ that can meaningfully constrain the role we propose the nervous system plays in our behaviour. The field of cognitive neuroscience has more fully come around to this view recently, with several high-profile papers arguing that good behavioural work is required to enable good neuroscientific work (e.g. Hasson et al., 2020; Jirsa et al., 2019; Krakauer et al., 2017). Krakauer et al explicitly cite the mechanism literature to note the need for a behavioural grounding scale theory to constrain neuroscience research and modelling. Ecological real parts and processes provide that grounding scale the way nothing else in psychology can right now (Golonka & Wilson, 2019a) putting us in the best position to develop a really effective cognitive neuroscience.

We as a field are ready to have these discussions, as evidenced by a recent issue of *Ecological Psychology* that presented a variety of contributions feeding into this new endeavour (de Wit & Withagen, 2019). In that issue, Golonka and Wilson (2019b) argued that ecological perceptual information about task dynamical properties identifies the correct behavioural grounding scale for understanding functional neural activity, and that the nervous system should therefore be considered in terms of how it brings that information and our effectivities together to form task-specific devices that successfully complement the task affordances. This was an explicitly mechanistic approach, building on Golonka and Wilson (2019a); information defines our grounding scale, and we took the real-components constraint as seriously as we could as we tried to derive the implications of interacting with real ecological information variables for understanding neural activity. We also then proposed a first draft analysis of how that neural activity might go on to serve as real components for other, more abstract, “representation-hungry” problems.

While it is early days and there is much work to be done on this framework, we found the mechanistic mindset provided us with an invaluable discipline as we developed our analysis; we couldn’t simply invoke any filler terms implemented in abstract representations⁸ to fill gaps. (This was important, because that is the kind of move we have criticised non-ecological neuroscience for making; Charles et al., 2014). The next

⁷Some ecological psychologists think we do not need to develop this scale for our explanations and I’d like to note here again that this move is not compulsory, from a mechanism point of view. Sometimes the explanation at the grounding scale is all that is required for a satisfactory explanation of the behaviour in question, and that’s just fine; it’s up to the scientists to decide if pursuing a multi-scale explanation is worthwhile. However, if they *do* want to pursue one, having the right grounding scale to constrain it is a definite requirement and the real components discipline will be invaluable to making that work as effective as possible.

⁸We did call the paper ‘Ecological Representations’, which in hindsight may have been a public relations error. Nevertheless, representations are invoked by cognitive science to solve particular problems, and the fact is that our ecological mechanistic analysis solves those problems in representational ways, just with more defensibly real parts and processes. It’s a tricky one, we know.

move is therefore to consider how best to dynamically characterise the informationally-constrained neural real parts that form one scale of a multi-scale behavioural mechanism. The special issue had some promising alternatives worth exploring, in the form of Transiently Assembled Local Neural Subsystems (TALoNS; Raja & Anderson, 2019), structured flow on manifolds (Jirsa et al., 2019), the formalisms of the Free Energy Principle (Bruineberg & Rietveld, 2019; see also Bruineberg et al., 2018; Bruineberg & Rietveld, 2014), and neural resonance⁹ (Fultot et al., 2019). Time will tell which approach will work best, but if the resulting work is properly grounded at the scale of ecological behavioural mechanisms then I suggest that our neuroscience will be a major leap forward compared to the functional analyses of non-ecological cognitive neuroscience.

Conclusions

The discipline that comes with mechanistic research and modelling leads to good scientific practice. It prevents us from panicking when at first we don't succeed; that just becomes a hint that we are grounding our explanation at the wrong scale. It stops us from filling gaps in our understanding with capacities such as 'efference copy' or 'temporal coding' and forces us to consider how to get that capacity only out of the real parts and processes our ontology provides. It provides strong reasons not to be reductionist, but multi-scale in our explanations, and guides how to do that carefully (Bechtel, 2009). And, of course, those explanations, implemented as dynamical systems models, are actually explanations, with all the benefits that come with that; we can use the model to study the system in ways we can't in the actual system, and it serves as a much more successful guide to discovery as we engage in iterative rounds of research and modelling. If a scientific research programme can be mechanistic, it should be.

Ecological psychology has argued for a long time that it cannot be mechanistic; our ontology doesn't permit it. We have therefore invested in dynamical and information-control research and modelling strategies instead. However, I have argued here that these are not as explanatory as ecological researchers have claimed, and the limitations can be seen in the empirical literature. If this is all we have access to, then so be it; we can still do perfectly useful science. I have argued instead that ecological psychology can be a fully mechanistic research and modelling programme, and that therefore we should work to become one, in order to gain all the explanatory benefits of the approach. I have done this by addressing the ecological objections to justify the claims that a) our ontology, while non-Cartesian, still provides allowable candidate real parts and processes, and b) we can avail ourselves of the decomposition and localisation research programme to study the systems these components form, using Bingham's research programme about coordinated rhythmic movement as a proof-of-concept. We can do it in theory, we can do it in practice, and the result is demonstrably better models and

⁹Maybe. My reading of this paper and the fact Turvey is one of the authors suggests to me that the offered account of neural resonance is to be used more like a dynamical covering law. However, a reviewer suggested and I agree that the resonance they discuss *is* tied to specific properties of groups of neurons, so, like the other options, neural resonance may end up being a useful dynamical characterisation of a real component in a future ecological neuroscience mechanistic model. I'll just flag this up here for people to consider, but don't read this as committing those authors to anything they may not agree with.

explanations. Uniquely in the cognitive sciences, the ecological programme of research and modelling can, and therefore I suggest should, explicitly work to become mechanistic.

ORCID

Andrew D. Wilson  <http://orcid.org/0000-0002-0459-0728>

References

- Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In *Psychology of Learning and Motivation* (Vol. 2, pp. 89–195). Academic Press.
- Baddeley, A. (2000). The episodic buffer: A new component of working memory? *Trends in Cognitive Sciences*, 4(11), 417–423. [https://doi.org/10.1016/S1364-6613\(00\)01538-2](https://doi.org/10.1016/S1364-6613(00)01538-2)
- Baddeley, A. D., & Hitch, G. (1974). Working memory. In *Psychology of learning and motivation* (Vol. 8, pp. 47–89). Academic press.
- Bechtel, W. (2009). Looking down, around, and up: Mechanistic explanation in psychology. *Philosophical Psychology*, 22(5), 543–564. <https://doi.org/10.1080/09515080903238948>
- Bechtel, W., & Abrahamsen, A. (2010). Dynamic mechanistic explanation: Computational modeling of circadian rhythms as an exemplar for cognitive science. *Studies in History and Philosophy of Science Part A*, 41(3), 321–333. <https://doi.org/10.1016/j.shpsa.2010.07.003>
- Bechtel, W., & Richardson, R. C. (2010). *Discovering complexity: Decomposition and localization as strategies in scientific research*. MIT press.
- Bingham, G. P. (1995). Dynamics and the problem of visual event recognition. In R. Port & T. van Gelder (Eds.), *Mind as motion: Dynamics, behavior and cognition* (pp. 403–448). MIT Press.
- Bingham, G. P. (2001). A perceptually driven dynamical model of rhythmic limb movement and bimanual coordination. *Proceedings of the 23rd Annual Conference of the Cognitive Science Society* (pp. 75–79). LEA Publishers.
- Bingham, G. P. (2004a). A perceptually driven dynamical model of bimanual rhythmic movement (and phase perception). *Ecological Psychology*, 16(1), 45–53. https://doi.org/10.1207/s15326969eco1601_6
- Bingham, G. P. (2004b). Another timing variable composed of state variables: Phase perception and phase driven oscillators. In *Advances in psychology* (Vol. 135, pp. 421–442). North-Holland.
- Bingham, G. P., Schmidt, R. C., & Zaal, F. (1999). Visual perception of the relative phasing of human limb movements. *Perception & Psychophysics*, 61(2), 246–258.
- Bingham, G. P., Zaal, F. T. J. M., Shull, J. A., & Collins, D. R. (2001). The effect of frequency on the visual perception of relative phase and phase variability of two oscillating objects. *Experimental Brain Research*, 136(4), 543–552. <https://doi.org/10.1007/s002210000610>
- Bogaerts, H., Buekers, M. J., Zaal, F. T., & Swinnen, S. P. (2003). When visuo-motor incongruence aids motor performance: The effect of perceiving motion structures during transformed visual feedback on bimanual coordination. *Behavioural Brain Research*, 138(1), 45–57.
- Bruineberg, J., Kiverstein, J., & Rietveld, E. (2018). The anticipating brain is not a scientist: The free-energy principle from an ecological-enactive perspective. *Synthese*, 195(6), 2417–2444. <https://doi.org/10.1007/s11229-016-1239-1>
- Bruineberg, J., & Rietveld, E. (2014). Self-organization, free energy minimization, and optimal grip on a field of affordances. *Frontiers in Human Neuroscience*, 8, 599.
- Bruineberg, J., & Rietveld, E. (2019). What's inside your head once you've figured out what your head's inside of. *Ecological Psychology*, 31(3), 198–217. <https://doi.org/10.1080/10407413.2019.1615204>

- Charles, E. P., Wilson, A. D., & Golonka, S. (2014). The most important thing neuropragmatism can do: Providing an alternative to 'cognitive' neuroscience. In J. Shook & T. Solymosi (Eds.), *Pragmatist neurophilosophy: American philosophy and the brain* (p. 127-149). Bloomsbury Academic.
- Chemero, A. (2009). *Radical embodied cognitive science*. MIT Press.
- Chemero, A., & Silberstein, M. (2008). After the philosophy of mind: Replacing scholasticism with science. *Philosophy of Science*, 75(1), 1–27. <https://doi.org/10.1086/587820>
- Craver, C. F. (2007). *Explaining the brain: Mechanisms and the mosaic unity of neuroscience*. Oxford University Press.
- Craver, C. F., & Darden, L. (2013). *In search of mechanisms: Discoveries across the life sciences*. University of Chicago Press.
- de Wit, M. M., & Withagen, R. (2019). What should a “Gibsonian neuroscience” look like? Introduction to the special issue. *Ecological Psychology*, 31(3), 147–151. <https://doi.org/10.1080/10407413.2019.1615203>
- Eronen, M. I. (2015). Levels of organization: A deflationary account. *Biology & Philosophy*, 30(1), 39–58. <https://doi.org/10.1007/s10539-014-9461-z>
- Fajen, B. R. (2007). Affordance-based control of visually guided action. *Ecological Psychology*, 19(4), 383–410. <https://doi.org/10.1080/10407410701557877>
- Feldman, A. G. (1986). Once more on the equilibrium-point hypothesis (lambda model) for motor control. *Journal of Motor Behavior*, 18(1), 17–54. <https://doi.org/10.1080/00222895.1986.10735369>
- Fontaine, R. J., Lee, T. D., & Swinnen, S. P. (1997). Learning a new bimanual coordination pattern: Reciprocal influences of intrinsic and to-be-learned patterns. *Canadian Journal of Experimental Psychology = Revue Canadienne de Psychologie Experimentale*, 51(1), 1–9.
- Frank, T. D., Daffertshofer, A., Peper, C. E., Beek, P. J., & Haken, H. (2000). Towards a comprehensive theory of brain activity:: Coupled oscillator systems under external forces. *Physica D: Nonlinear Phenomena*, 144(1–2), 62–86. [https://doi.org/10.1016/S0167-2789\(00\)00071-3](https://doi.org/10.1016/S0167-2789(00)00071-3)
- Fultot, M., Adrian Frazier, P., Turvey, M. T., & Carello, C. (2019). What are nervous systems for? *Ecological Psychology*, 31(3), 218–234. <https://doi.org/10.1080/10407413.2019.1615205>
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Houghton Mifflin.
- Goldfield, E. C., Kay, B. A., & Warren, W. H. Jr, (1993). Infant bouncing: The assembly and tuning of action systems. *Child Development*, 64(4), 1128–1142.
- Golonka, S., & Wilson, A. D. (2019a). Ecological mechanisms in cognitive science. *Theory & Psychology*, 29(5), 676–696. <https://doi.org/10.1177/0959354319877686>
- Golonka, S., & Wilson, A. D. (2019b). Ecological representations. *Ecological Psychology*, 31(3), 235–253. <https://doi.org/10.1080/10407413.2019.1615224>
- Hasson, U., Nastase, S. A., & Goldstein, A. (2020). Direct fit to nature: An evolutionary perspective on biological and artificial neural networks. *Neuron*, 105(3), 416–434. <https://doi.org/10.1016/j.neuron.2019.12.002>
- Hatsopoulos, N. G., & Warren, W. H. Jr, (1996). Resonance tuning in rhythmic arm movements. *Journal of Motor Behavior*, 28(1), 3–14.
- Herth, R. A., Zhu, Q., & Bingham, G. P. (2021). The role of intentionality in the performance of a learned 90° bimanual rhythmic coordination during frequency scaling: Data and model. *Experimental Brain Research*, 239(10), 3059–3075.
- Hitch, G. J., & Baddeley, A. D. (1976). Verbal reasoning and working memory. *Quarterly Journal of Experimental Psychology*, 28(4), 603–621. <https://doi.org/10.1080/14640747608400587>
- Jirsa, V. K., McIntosh, A. R., & Huys, R. (2019). Grand Unified Theories of the brain need better understanding of behavior: The two-tiered emergence of function. *Ecological Psychology*, 31(3), 152–165. <https://doi.org/10.1080/10407413.2019.1615207>
- Kaplan, D. M., & Bechtel, W. (2011). Dynamical models: An alternative or complement to mechanistic explanations? *Topics in Cognitive Science*, 3(2), 438–444. <https://doi.org/10.1111/j.1756-8765.2011.01147.x>

- Kaplan, D. M., & Craver, C. F. (2011). The explanatory force of dynamical and mathematical models in neuroscience: A mechanistic perspective. *Philosophy of Science*, 78(4), 601–627. <https://doi.org/10.1086/661755>
- Kay, B. A., Kelso, J. A., Saltzman, E. L., & Schönner, G. (1987). Space–time behavior of single and bimanual rhythmical movements: Data and limit cycle model. *Journal of Experimental Psychology. Human Perception and Performance*, 13(2), 178–192.
- Kay, B. A., Saltzman, E. L., & Kelso, J. A. (1991). Steady-state and perturbed rhythmical movements: A dynamical analysis. *Journal of Experimental Psychology. Human Perception and Performance*, 17(1), 183–197.
- Kelso, J. S. (1995). *Dynamic patterns: The self-organization of brain and behavior*. MIT press.
- Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., MacIver, M. A., & Poeppel, D. (2017). Neuroscience needs behavior: Correcting a reductionist bias. *Neuron*, 93(3), 480–490.
- Leach, D., Kolokotroni, Z., & Wilson, A. D. (2021a). Perceptual information supports transfer of learning in coordinated rhythmic movement. *Psychological Research*, 85(3), 1167–1182.
- Leach, D., Kolokotroni, Z., & Wilson, A. D. (2021b). The ecological task dynamics of learning and transfer in coordinated rhythmic movement. *Frontiers in Human Neuroscience*, 506
- Menary, R. (2007). *Cognitive integration: Mind and cognition unbounded*. Springer.
- Milkowski, M., Hohol, M., & Nowakowski, P. (2019). Mechanisms in psychology: The road towards unity? *Theory & Psychology*, 29(5), 567–578. <https://doi.org/10.1177/0959354319875218>
- Miller, G. A. (1956). The magic number seven plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63(2), 81–97. <https://doi.org/10.1037/h0043158>
- Nam, H., Goldstein, L., & Saltzman, E. (2009). Self-organization of syllable structure: A coupled oscillator model. In Pellegrino, F., Marsico, E., Chitoran, I., & Coupé, C. (Eds.), *Approaches to phonological complexity* (pp. 297–328). De Gruyter Mouton.
- Nordham, C. A., Tognoli, E., Fuchs, A., & Kelso, J. S. (2018). How interpersonal coordination affects individual behavior (and vice versa): Experimental analysis and adaptive HKB model of social memory. *Ecological Psychology : A Publication of the International Society for Ecological Psychology*, 30(3), 224–249.
- Piccinini, G. (2020). *Neurocognitive mechanisms: Explaining biological cognition*. Oxford University Press.
- Pinker, S., & Prince, A. (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition*, 28(1–2), 73–193. [https://doi.org/10.1016/0010-0277\(88\)90032-7](https://doi.org/10.1016/0010-0277(88)90032-7)
- Port, R. F., & Kaipainen, M. (2002). Temporal attractors in timing: Applying the HKB model to speech. *IULC Working Papers*, 2(2)
- Raja, V., & Anderson, M. L. (2019). Radical embodied cognitive neuroscience. *Ecological Psychology*, 31(3), 166–181. <https://doi.org/10.1080/10407413.2019.1615213>
- Rumelhart, D. E., & McClelland, J. L. (1986). On learning the past tenses of English verbs. In J. L. McClelland & D. E. Rumelhart (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition, Vol. 2. Psychological and biological models* (pp. 216–271). MIT Press.
- Schmidt, R. C., Carello, C., & Turvey, M. T. (1990). Phase transitions and critical fluctuations in the visual coordination of rhythmic movements between people. *Journal of Experimental Psychology. Human Perception and Performance*, 16(2), 227–247.
- Schmidt, R. C., & Richardson, M. J. (2007). Dynamics of interpersonal coordination. In Fuchs, A., & Jirsa, V. K. (Eds.), *Coordination: Neural, behavioral and social dynamics* (pp. 281–308). Springer Science & Business Media.
- Silberstein, M. (2021). Constraints on localization and decomposition as explanatory strategies in the biological sciences 2.0. In Calzavarini F., & Viola M. (Eds.), *Neural Mechanisms. Studies in Brain and Mind* (Vol. 17). Cham: Springer. https://doi.org/10.1007/978-3-030-54092-0_16
- Silberstein, M., & Chemero, A. (2013). Constraints on localization and decomposition as explanatory strategies in the biological sciences. *Philosophy of Science*, 80(5), 958–970. <https://doi.org/10.1086/674533>

- Snapp-Childs, W., Wilson, A. D., & Bingham, G. P. (2011). The stability of rhythmic movement coordination depends on relative speed: The Bingham model supported. *Experimental Brain Research*, 215(2), 89–100.
- Snapp-Childs, W., Wilson, A. D., & Bingham, G. P. (2015). Transfer of learning between unimanual and bimanual rhythmic movement coordination: Transfer is a function of the task dynamic. *Experimental Brain Research*, 233(7), 2225–2238.
- Stepp, N., Chemero, A., & Turvey, M. T. (2011). Philosophy for the rest of cognitive science. *Topics in Cognitive Science*, 3(2), 425–437.
- Tognoli, E., Zhang, M., Fuchs, A., Beetle, C., & Kelso, J. S. (2020). Coordination Dynamics: A foundation for understanding social behavior. *Frontiers in Human Neuroscience*, 317. <https://doi.org/10.3389/fnhum.2020.00317>
- Tulving, E. (1972). Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), *Organization of Memory* (pp. 381–403). Academic.
- Turvey, M. T. (2019). *Lectures on perception: An ecological perspective*. Routledge.
- Warren, W. H. (2006). The dynamics of perception and action. *Psychological Review*, 113(2), 358–389. <https://doi.org/10.1037/0033-295X.113.2.358>
- Wenderoth, N., Bock, O., & Krohn, R. (2002). Learning a new bimanual coordination pattern is influenced by existing attractors. *Motor Control*, 6(2), 166–182.
- Wilson, A. D., & Bingham, G. P. (2008). Identifying the information for the visual perception of relative phase. *Perception & Psychophysics*, 70(3), 465–476. <https://doi.org/10.3758/pp.70.3.465>
- Wilson, A., Bingham, G. P., & Craig, J. C. (2003). Haptic perception of phase variability. *Journal of Experimental Psychology: Human Perception and Performance*, 29, 1179–1190.
- Wilson, A. D., Snapp-Childs, W., & Bingham, G. P. (2010). Perceptual learning immediately yields new stable motor coordination. *Journal of Experimental Psychology: Human Perception and Performance*, 36(6), 1508–1514.
- Wimmers, R. H., Beek, P. J., & van Wieringen, P. C. (1992). Phase transitions in rhythmic tracking movements: A case of unilateral coupling. *Human Movement Science*, 11(1–2), 217–226. [https://doi.org/10.1016/0167-9457\(92\)90062-G](https://doi.org/10.1016/0167-9457(92)90062-G)
- Zaal, F., Bingham, G. P., & Schmidt, R. C. (2000). Visual perception of relative phase and phase variability. *Journal of Experimental Psychology: Human Perception and Performance*, 26(3), 1209–1220.
- Zanone, P. G., & Kelso, J. A. S. (1994). The coordination dynamics of learning: Theoretical structure and experimental agenda. In Swinnen, S. P., Massion, J., Heuer, H., & Casaer, P. (Eds.), *Interlimb coordination: Neural, dynamical, and cognitive constraints*. Academic Press.