Ecological Representations

Sabrina Golonka*

&

Andrew D Wilson

*Corresponding Author

Psychology, School of Social Sciences

Leeds Beckett University

Leeds, LH1 3HE, UK

Email: s.golonka@leedsbeckett.ac.uk / psychscientists@gmail.com

Web: https://cognitioninaction.wordpress.com/

Twitter: @PsychScientists

Abstract

Cognitive science has three main motivations for claiming that cognition requires representation; the need for intentional access to the world, poverty of perceptual access to that world, and the need to support 'higher-order' cognition. In addition to these motivations, all representational systems must address two major problems: symbol grounding and system-detectable error. Here we argue that James J Gibson's ecological information fits the basic definition of a representation, solves both problems and immediately addresses the first two motivations. We then develop an argument (begun in Golonka, 2015) that *informational representations* and the resulting *neural representations* can also support 'higher-order' cognition and provides an ecological framework for understanding neural activity. Ecological psychology can be a complete theory of cognition, and the key is the way that information represents the world.

Keywords: psychology, ecological information, knowledge representation, symbol grounding, intentionality

Introduction

Since the cognitive revolution, theories of cognition have critically involved the notion of *representation*. There are three non-trivial motivations for this notion, rooted in analyses of the kind of thing a cognitive system must be. Specifically, cognition is a 1) flexible and intentional system that is implemented in a physical system, which 2) interacts with impoverished sensory signals, and 3) can do things such as think about things in their absence. All these seem to demand representational support, and the current mainstream consensus is that cognition uses *mental* representations that are *internal* to the system (i.e. located in the brain) and *computational* in nature.

Mental representations are not without their problems, however. The two major and, as yet, unsolved problems are the issues of *symbol grounding* (Harnad, 1990) and *system-detectable error* (Bickhard, 2009). These problems, and a growing recognition that perception is in fact rich and informative, have motivated 'radical' theories of embodied cognition that argue representations are not required to support cognition (e.g. Chemero, 2009). However, the major successes of non-representational accounts have been restricted to perception-action tasks and they have yet to develop any widely accepted explanations for the 'high-order' cognitive activities driving motivation (3), or the activity of the brain. As such, they are often ruled out as having the potential to offer a complete account of cognition.

Human cognition has clearly found a way to cope with motivations (1) and (3); we are biologically-implemented intentional cognitive systems who can think about things in their absence. Poverty of stimulus problems are therefore either being solved or (as ecological psychologists argue) are not a problem. This leaves us with the following possibilities for what kind of thing cognition is:

1) human cognition is representational and the implementation of these representations solves the two major problems is an, as yet, unknown way (the approach we will take here)

2) human cognition is representational, but we are mistaken about the two problems and don't need to worry about them in our theories (the approach adopted by most of modern cognitive science[1]), or
3) human cognition is not representational and we are mistaken that motivations (1) and (3) require representations (the radical approach described above)

Right now, the field is at an impasse. Radical approaches are doing well in explaining real-time skilled engagement with the environment, but there's no widely accepted non-representational account of higher-order cognition or neural activity. Some argue for a blend of approaches in which representations kick in at these higher levels, although given how very different in kind these approaches are (Turvey et al, 1981; Wilson, 2018) it's not clear how to make this work.

We believe there might be a way to reframe the overall debate in a way that preserves the good work of all sides and provides a path for a more comprehensive ecological research program. We will agree that cognition requires representations. However, we propose that Gibsonian ecological information perfectly fits the definition of a representation. These *informational* representations solve both the symbol grounding and system-detectable error problems, and they constrain the form (and empirical investigation) of *neural* representations caused by interacting with information. These two *ecological representations* then address all three motivations for representational accounts described above, including, as we develop below, the major challenge of supporting 'higher-order' cognition.

The rest of this paper lays this argument out in full. First, we review the three motivations for representations in more detail. We lay out the reasons why cognitive science settled on computational mental representations as the solution, and review the two problems that came with that solution. We will then discuss our approach to understanding what a representation is, identify

[1] Although see Wilson, 2018 for an analysis of how this issue has recently resurfaced in the interface theory of perception (Hoffman et al, 2015).

that ecological information fits the bill and solves the problems, and show how this provides a framework for interpreting neural activity ecologically. Information quite naturally addresses motivations (1) and (2); we will then do the critical work of showing how it can address motivation (3).

From Mental to Ecological Representations

Representations: Three Motivations, a Solution, Two Problems

A good theory should explain the characteristics of a set of related phenomena. Newell (1980) proposed a list of characteristics that a theory of cognition needs to eventually address. The most important is universality, which is the idea that a cognitive system must be able to behave as an (almost) arbitrary function of the environment. It must be able to be 'about' anything it encounters in the world; cognition must be *intentional*. Newell also emphasized that implementing such a flexible, adaptive system in an actual, physical system was a major problem (Bechtel, 1998; Stich, 1992).

Early cognitive scientists (e.g. Fodor, 1980; Pylyshyn, 1989) argued that universality could only be implemented in a *computational* system; these have the necessary indefinite flexibility and can be realized in a variety of physical systems. They then noted that computational systems are necessarily representational[2]. The primary motivation for treating cognition as necessarily representational is therefore *the need for intentionality*.

We can divide the intentionality requirement into two sub-cases; motivations (2) and (3). First, we can act intentionally towards properties of objects/events that aren't in direct mechanical contact with us, via our perceptual contact with those objects/events. That contact is impoverished and ambiguous, however, and so successful intentional interaction with the environment requires the computational representational systems to enrich that contact with stored knowledge. Second, we

[2] But see Picninini 2004, 2008 for an argument against this

can act intentionally towards higher-order properties of objects/events that might not be present in, or might even be in opposition to, the current environment. Here, our perceptual contact with such entities is non-existent, so again representations are required to suitably enrich our experience to support successful intentional interactions.

These cases require representations to fill *gaps* in the causal chain of events leading from world to functional behavior in that world. If we are sitting in a tree and the branch we are sitting on breaks, the force of gravity causes us to be displaced and there's no gap in the chain of cause and effect. However, if we move because a branch falls nearby, or because it is windy and we are worried that a branch might fall, a gap opens up in the causal chain of events. There was no mechanical linkage between the accelerating tree branch and our body; in the latter case, there wasn't even an accelerating tree branch at all, just the worry of one. Yet, we moved. Obviously, we don't move by accident or by magic. There are reasons and explanations, but these reside in a psychological level of explanation, which is why we need a psychological entity – like representations – to fill the gap.

This analysis makes a strong case for computational mental representations as key players in our theories of cognition, because only these systems have the necessary properties to address all three motivations (Fodor, 1980; Newell, 1980; Pylyshyn, 1989).

However, these kinds of representations are not a slam-dunk solution. If you have 20 years to spare, we encourage you to peruse this literature and note the many ongoing debates about how to exactly implement these representations, and the various consequences of the methods. When you are done, we hope you agree with our modest assessment that there are certain important concerns with mental representations that have not been laid to rest.

The first issue is that symbolic/representational systems do not necessarily come with intrinsic access to intentional content and, thus, cannot actually be 'about' anything without additional help.

The old-school name for this is the *symbol grounding problem*[3] (Bickhard, 2009; Harnad, 1990; Searle, 1980). Symbolic representations suffer from the arbitrary mapping between the structure of the representation and the thing it's meant to represent, so attempts to solve the symbol grounding problem try to link abstract or higher-order mental representations to more basic perceptual representations (e.g. Barsalou, 1999; Harnad, 1990). However, standard cognitive psychology treats perception extensionally (Turvey et al, 1981; Wilson, 2018) which means even the lower perceptual representations lack internally defined content and there is no grounding for the higher processes to inherit. This "infinite regress of interpreters interpreting" (Bickhard, 2009, p 573) is endemic to any representational account where content is defined externally, which Bickhard argues includes Milikan, Fodor, Dretske, and Cummins. Therefore, none of these mental representations are grounded and their contents are remain inaccessible to the organism. So despite falling out of the conversation somewhat, symbol grounding remains an unsolved and serious problem for mental representations.

The second, related problem is that symbolic/representational systems have no frame of reference to identify when they are making errors and thus cannot adapt their behavior to become better attuned to their environments (the old-school name is the *system-detectable error* problem; Bickhard, 2009). Error-guided learning requires that a system is able to detect representational errors. The ability to detect representational errors requires access to representational content, which cannot happen if representations are externally defined and not grounded.

We now have a job description for a cognitive theory. The proposed components must support intentional behavior, including cases where the behaviorally-relevant properties of objects or events are spatially and/or temporally distal. Furthermore, they must have grounded intentional content and enable system-detectable error. To date, neither computational mental representational nor

[3] Not everyone thinks the problem remains (Steels, 1997; Taddeo & Floridi, 2007). But the debate remains (Bringsjord, 2015) and so, for the people who, like us, still worry about this problem, we endeavour to solve it.

anti-representational approaches have been able to fulfil this entire job description. Here we argue that two flavors of *ecological representation* (*informational* representations and the *neural* representations they create) fill the full description extremely well, up to and including the daunting question of 'higher-order cognition'.

Representations: Not a Definition

To develop this argument, we need to take a step back and consider just what the essence of a cognitive representation is. This is a complicated topic, to say the least. We are not philosophers and cannot here bring the same depth of coverage to the historical place of representations in cognitive science as a philosopher would. Luckily for us, a philosopher (Ramsey, 2007) already has, so we will develop our argument from there. In addition, Cummins (1989) suggests that trying to pin down a conception of representation is pointless without connecting it to the theory within which it is supposed to operate. This is the other reason we rely on Ramsey; he does a fantastic job of keeping the philosophical nature of representation tethered to the empirical goals of cognitive science. This is a long preamble to simply say that "If what [we] say does not accurately describe your own favorite theory or model, [we] ask that you consider [our] claims in light of what you know about more general conventions, attitudes, assumptions and traditions."

For the concept of representation to have any bite, it has to mean something more than simply an "inner" or "causally relevant" state. (Ramsay, 2007 p 8). Ramsey's hard work points to the fact that regardless of format or other bells and whistles, all cognitive theories of representation rest upon "the basic idea that inner structures in some way serve to stand for, *designate*, or mean something else" (Ramsay, p 3, emphasis ours). Newell (1980) usefully defines designation like this:

> Designation: An entity X designates an entity Y relative to a process P, if, when P takes X as input, its behavior depends on Y.
>
> There are two keys to this definition: First, the concept is grounded in the behavior of a process. Thus, the implications of designation will depend on the nature of this

> process. Second, there is action at a distance . . . This is the symbolic aspect, that having X (the symbol) is tantamount to having Y (the thing designated) for the purposes of process P (Newell, 1980, p. 156).

The next step in the argument is to demonstrate that the function of ecological information is to designate the world. We will show this with a careful consideration of the nature of this information; how it comes to be, and how it is used by an organism.

Ecological Information is a Representation

Ecological information is constituted by higher-order relational patterns in energy arrays. These patterns are created by the lawful interaction of the energy with the dynamics of the world (Turvey et al, 1981) and are used by organisms to perceive that world (Bingham, 1988; Gibson, 1966, 1979; Michaels & Carello, 1981). As we will now explain, the most important feature of information is that while it can specify world properties, it is not, in fact cannot be, identical to those properties. Because organisms are only in perceptual contact with information and not the property, their behavior is shaped by the form of the information and not the form of the property. The function of information is to stand in for and designate the dynamical world, and is thus usefully considered as a representation.

The Kinematic Specification of the Dynamical World

Properties of objects and events in the world change over space and time in ways that reflect the composition and organization of these properties. These properties and changes can only be comprehensively described at the level of *dynamics* (Bingham, 1988, 1995) which allows units of time, position (and its derivatives), and mass. Successfully interacting with a task therefore requires that the behavior *complements* the dynamics of the task, and intentional systems (as described earlier) are therefore systems whose behavior is appropriately organized with respect to task dynamics.

There is a problem, however. The particular type of intentional systems that humans and other animals are do not have unmediated access to task dynamics. This is because almost all of the behaviorally relevant dynamics in the world are 'over there'; they are spatially or temporally distal and not in mechanical contact with the organism. Fortunately, those distant dynamics are embedded in the various ambient energy media in our ecology, and the spatial-temporal structure in the dynamics *projects* into these adjacent energy arrays. We are also embedded in these media, and so this projection can connect task dynamics to organism.

There is a new problem, however. These projections are not identical to the underlying dynamical properties. Specifically, they are kinematic projections (Bingham, 1988; Turvey et al, 1981). Kinematics is a level of description that only refers to motions; the units are time, position (and its derivatives), but *not* mass. The projection has lost a critical dimension, and because of this 'perceptual bottleneck' (Bingham, 1988) kinematic information cannot be identical to the dynamical world.

Fortunately, the way dynamics are projected into kinematic information is a law-based process. The details of the projection are unambiguously related to the details of the dynamic that caused it, and different dynamics project differently. Consequently, the kinematic patterns can *specify* (i.e., map 1:1 to) the dynamics that caused them, without needing to be identical to them (Gibson, 1979; Runeson & Frykholm, 1983; Turvey et al, 1981).

Ecological Information Designates & Therefore Represents the Dynamical World

Organisms need access to task dynamics but are only in contact with kinematic projections of those dynamics. These projections specify the dynamics, and so can be used as information for those dynamics. We can now see that the function of the information for the organism is to represent the world. With reference to Newell's (1980) definition described above, when behavioral process P takes information variable X as input, P's behavior depends on the specified dynamical property Y; there is the required 'action-at-a-distance' and the structure of P is only explainable with reference

to X and its relationship to Y. Therefore, the consequence of ecological information on perception-action systems is best explained by invoking the ability of information to designate, and therefore represent, properties of objects and events in the world. In fact, we'd struggle to explain why animals react to structure in kinematic arrays without invoking the ability of these patterns to stand in for world properties.

Neural Activity Designates and Therefore Represents Information

Informational representations provide a mechanism connecting the world to perceptual systems. We also need a mechanism connecting informational representations to behavior. This is going to involve the nervous system, but our ecological analysis presents a decidedly novel way to characterize what the nervous system needs to accomplish. Specifically, we must explain how informational representations cause changes in nervous system activity such that resultant actions correlate with and/or complement environmental properties.

The general finding in ecological behavioral research shows that the form of a behavior maps tightly onto the structure of the information involved. For example, the behavioral level characteristics of coordinated rhythmic movement are accounted for by the structure of relative direction, the informational representation of the world property relative phase (Bingham, 2001, 2004a, b, reviewed in Golonka & Wilson, 2012; see Wilson et al, 2005a for an empirical demonstration). A parsimonious hypothesis is that the nervous system enables this correlation by preserving or systematically transforming ecological informational structure.

Explicitly ecological analyses of neural activity (i.e. tracking the neural consequences of interacting with ecological information) are currently rare, but the few existing studies support our hypothesis (e.g. Agyei et al, 2015; Magrassi et al (2015); van der Meer et al, 2012; van der Weel & van der Meer, 2009). For example, van der Meer et al measured the magnitude of the correlation between an information variable and neural activity. They found that the strength of this relationship predicted task performance. In fact, changes in the magnitude of the correlation in an individual predicts

changes in performance over time. In other words, task performance (the extent to which behavior complements task dynamics) depends on how well informational structure is preserved in neural activity.

This reconceptualization of the neural code follows naturally from the ecological analysis of behavior. Rather than operating by synchronizing oscillations at various frequencies, the language of the brain may be more akin to the continuously unfolding nature of informational representations. The existence of this ongoing activity, so-called travelling waves, has been known about for decades (see Hughes, 1995 for a review), but it has not make much of an impact on neuroscience. One reason is that common neuropsychological techniques (e.g., EEG recording at the scalp, averaging across trials) actually conspire to obscure travelling-wave effects (Muller et al, 2018). In addition, there is no way to really understand what travelling waves are doing without looking at the structure of energy entering the system, and no one has yet added Gibson to the mix.

Neural activity therefore closes the gap from information to action by designating that information to the action systems. Therefore, we think it likely that (at least some of) the neural activity caused by informational representations will qualify as a neural representation of that information, which can now be interpreted with reference to the structure of the information.

Are these neural representations just those familiar from mainstream cognitive neuroscience? We think not. Any psychological power possessed by these neural representations comes from their relationship to external, informational representations. Their job would not be to enrich, model or predict anything about that information. This means that understanding the function and structure of neural representations requires understanding the structure and environmental cause of ecological information, which is not how cognitive neuroscience currently guides its work. Second, we will argue below that there are limits on the ability to call upon neural representations in the absence of the corresponding ecological information, something mental representations are

explicitly cited as being able to do. These neural representations are, therefore, not implementing the mental representations of the standard cognitive approach.

We now turn to the benefits of adopting these *ecological representations*.

Ecological Information Solves the Symbol Grounding Problem

The symbol grounding problem is eliminated if the relationship between the world and the medium onto which it is projected is entirely causal (meaning, non-arbitrary, lawful) because the higher level systems inherit the lower level grounding (Harnad, 1990). Informational representations are inarguably grounded. The link between the structure of an informational representation and the property of the world it represents is not arbitrary. Information is the result of a lawful process that projects dynamics into kinematics (Turvey et al, 1981) and the lawfulness enables those kinematics to be specific to the dynamics they represent (Runeson & Frykholm, 1983). Not being identical makes information a representation; specificity makes information an effective representation; the laws make information a grounded representation. This solution is then allowed to propagate up: neural representations caused by informational representations inherit that grounding, as does the resulting behavior.

Ecological Information Supports System-Detectable Error

System detectable error requires a representation to be internally related to its contents. Ecological information (the representational vehicle) is internally related to events in the world (the representational content); nothing additional to the information is required for it to mean what it does and be used to coordinate and control behavior. Errors occur when the information variable being used does not specify the currently relevant task dynamics (because it is the wrong variable to use, or because of an experimental perturbation). This mismatch is readily available to the organism and can drive corrections.

There is a growing experimental literature on this type of information switching (e.g. Jacobs, Michaels & Runeson, 2000; Jacobs, Runeson & Michaels, 2001; van der Meer et al, 2012; Withagen

& van Wermeskerken, 2009) and this has led to the development of the direct learning theory (Jacobs & Michaels, 2007). As an example, the relative phase of a coordinated rhythmic movement is typically perceived using relative direction of motion as the information (Bingham, 2001, 2004a, b; Wilson & Bingham, 2008; Wilson et al, 2005b). Relative direction is stable and clearly detected at 0° and 180° and so these coordination behaviors are stable as well (e.g. Wilson et al, 2005a). Relative direction is maximally variable at 90° meaning that relative direction information cannot support stable, functional coordination behavior at 90°. However, people can quickly learn to produce coordinated behavior at 90° (e.g. Wilson, Snapp-Childs & Bingham, 2010; Wilson, Snapp-Childs, Coats & Bingham, 2010) and they do so by switching to a new, more stable information variable (relative position; Wilson & Bingham, 2008). In our representational terminology, relative direction is a representation of relative phase that enables stable behavior at 0° and 180° but not 90°. This instability at 90° is detected and drives perceptual learning of a new representation, relative position, which now enables stable behavior at 90°. Information supports system-detectable error.

Ecological Information and the Motivations for Representations

Motivation 1 - the need to add intentionality: Because informational representations specify biologically relevant task dynamical properties (rather than perceptual primitives), they are inherently meaningful (Turvey et al, 1981) and also support our behavioral flexibility (Newell's universality). Even a completely novel object or event will have *properties* in common with things we have already encountered, which will cause familiar information and provide a basis for exploring the new context via error-driven learning made possible by grounding. Thus, intentional behavior follows from coordination with informational representations.

Motivation 2 - poverty of stimulus: This is the easy one to address; there is no poverty of stimulus. The lawful process by which the dynamical world is projected into information enables the kinematic specification to those dynamics. While there is a physical distance between us and most of the

world, this distance is literally filled with structure in energy media that is specific to (and represents) biologically relevant properties of distal objects and events.

Summary

The paper so far has taken what we already know about ecological information and identified that a) it fits the definition of a representation as an entity that designates, and b) that information is the only kind of behaviorally relevant representation that can fulfill the cognitive theory job description without falling victim to the serious problems that representations must address. The rest of this paper takes this analysis as a starting point and argues informational representations, combined with the neural representations they create, can support 'higher order cognition' and 'thinking about things in their absence'.

Thinking Ecologically About Things in Their Absence

How do we explain how behavior is shaped by things that are not perceptually specified right now? This is, to put it mildly, the issue ecological accounts have yet to tackle in detail. Research on ecological information focuses on the continuous control of action, which has told us a great deal about how information structures our here-and-now, real-time, online interactions with the world. But even when people are happy to cede ground to the ecological approach on the basis of these results, they invariably point out that at least some of our behavior does not involve interacting with things that are present and creating information (Clark & Toribio, 1994). Even Barsalou (1999) hypothesized that those experiences are later reified into perceptual symbols which could form the basis of "higher" cognitive functioning.

Action Control, Action Selection, Neural Representations

Our developing solution begins by identifying that information can not only control actions; it can also select them (Golonka, 2015). Action selection occurs when an organism chooses between alternatives, changes from one task to another, or parameterizes the performance of the current

task. When a friend verbally tells you to 'pick up the red cup', information in the auditory signal enables you to select which of two cups you pick up, although you then need to use visual and proprioceptive information (specifying the location of the cup and the movement of your arm) to actually implement and control the action.

The two roles (action control and action selection) place different requirements on information. In order for information to support action *control*, it must change in behaviorally-relevant real time as a direct function of some task-relevant property in the environment. In other words, it must continuously specify the current state of the world that created the information (e.g. as relative direction does for relative phase; Wilson & Bingham, 2008). This is what enables informational representations to support real-time coupling of behavior to properties currently present in the task environment. In order for information to support action *selection*, the task-relevant property might not even be present in the local task space and the information variable's structure does not need to relate in a lawful way to this property. If one encounters a door that says "Danger: Bear Inside," the task relevant properties (specifying the existence of a bear) do not structure patterns in ambient light that reach your retina because the bear is occluded by a door. However, ecological information caused by the sign on the door can cause neural activity that participates in selecting actions related to these distal properties (e.g., avoidance). From the first person perspective of the organism, it is just interacting with information. But while all ecological information is lawfully related to the properties of objects or events in the world that create the information, organisms are not law-bound to use that information in a particular way. Following Golonka (2015), when the behavioral consequences of information are not related to the object or event that caused the information, we say information has had a *conventional* (as opposed to law-based) effect on behavior.[4]

[4] This notion of conventionality has much in common with Pattee's ideas about symbols in biological systems (e.g., Pattee, 1972).

For our purposes, the relevant distinction between information used to select versus control actions concerns the stability of the neural activity caused by different information variables, especially when the corresponding information is not currently present in the environment. There is no convincing evidence that we can instantiate a neural representation of information sufficient to support action control unless the relevant information is present in the current environment (or was present recently enough to calibrate activity). One example of this that experienced drivers, despite being able to successfully steer a real car, are unable to realistically mime the action of steering, and often do so in ways that would have catastrophic consequences in actual driving contexts (Wallis, et al, 2007). Knowing how to steer in a real driving context is not the same as, and does not entail, being able to instantiate a neural representation of steering absent that context. This basic distinction (in the form of the online vs offline control of actions) also shows up in most of the core work around the two visual streams hypothesis; see Goodale, et al, (2004) for a review.

In contrast, we are often able to instantiate a neural representation of information used in action selection if 1) we have an appropriate precipitating event and 2) the structure of the information is simple, short, and/or well-practiced and stereotypical enough to have had a reliable functional effect on corresponding neural activity during learning. In humans, a familiar example of such neural representations is our ability to use inner speech. The structure of individual words for an experienced language user is simple, short, well-practiced and relatively stereotypical. The right precipitating event (e.g., reading the sign “Danger: Bear Inside”) can reliably instantiate a neural representation of the acoustic information caused by pronouncing these words. The result is that we “hear” the words in our heads (e.g. Breen & Clifton, 2013). This example relies on a close relationship between information present in the moment and the representation (i.e., they contain the same words in different modalities), but this connection is not obligatory. We could easily imagine training someone on a convention that a red circle on a door means that there is a bear inside. In this case, the information created by the colored circle causes neural activity related to the auditory information caused by the word “bear.” This neural activity functions as a neural

representation of the acoustic ecological information for the word "bear."[5] The informational representation of the word "bear" can impact action selection – for example, by selecting neural representations of words like "run" and "away." Once a person is trained on this convention, information for the red circle may also directly impact selection of actions related to avoidance. The actual escape from the bear, however, will require access to online information suitable for action control; the neural representation of the word "bear" cannot tell your legs how to move with respect to the supporting surface of the floor.

This is a simple example of how ecological information can enable functional behavior with respect to things not in the present environment, via things that are. There are two things worth drawing attention to. First, the hypothetical neural representations related to ecological information only have the power to select actions, not to control them. Second, these hypothetical neural representations aren't simply mental representations grounded in perceptual experience. They are re-instantiations of neural activity caused by ecological information. Such re-instantiations can have a phenomenality; it feels like something to hear language in our heads or to imagine someone's face; and they can have a consequence on future behavior by impacting selection of actions or selection of other neural representations (e.g., we continue our train of inner speech). But these neural representations, while internal, are not the mental representations of standard cognitive theories.

Ecological Higher Order Cognition

With this analysis in place, we will now discuss how neural representations of action-selecting information could enable some of the trickier aspects of human cognition. First, let us recall that we are not equating neural representations with all neural activity. Only neural activity that meets our working definition of a representation (i.e., that designates with respect to corresponding ecological information) will be considered a neural representation. Second, let us recall that not all neural

[5] Another example is how trained language users experience the heart symbol in the famous 'I ❤ NY' as if it were the spoken word "heart."

representations will be stable enough to instantiate in the absence of the corresponding informational representation. In fact, there should be a distribution of stability, such that some neural representations can be re-instantiated quite accurately (in the sense of a strong systematic relationship between neural activity and information, rather than an exact replica of the structure of the information), some with a certain degree of accuracy, some with very poor accuracy, and some, not at all. For the purposes of this discussion, we are concerned with the subset of neural activity that can be considered a representation and that can be re-instantiated with fair to good accuracy.

Barsalou and others all assume that knowledge systems must be both *conceptual* and *componential*, to allow complex expressions to be decomposed and new expressions to be built up, i.e. *productivity* (Chomsky, 1957; Dietrich & Markman, 2003; Fodor & Pylyshyn, 1988; Haugeland, 1991). Conceptual systems are removed from the particulars of a situation – they can represent general cases (concepts) rather than individuals (Barsalou, 1999). Componential systems contain parts that can be combined (and re-combined) according to, e.g., a recursive syntax. These features are realized in symbols systems, which is what makes them so good at supporting counterfactual thinking and context-dependent flexibility[6]. To be clear, the stipulation that knowledge systems must be conceptual and componential is so that knowledge systems can support counterfactual thinking, etc. To our minds, this means that the main challenge is to show that an ecological approach can support counterfactual thinking, etc., whether or not this solution involves a conceptual and componential system. However, as it turns out, we think that certain features of the two ecological representations mean that they enable conceptual, componential systems.

It is uncontroversial to say that developing a *concept* requires experience with multiple individuals of a type. For us, this would involve (at a minimum) repeated exposure to ecological information variables specifying properties of a given type. The neural activity caused by this repeated exposure

[6] S*ystematicity* is also discussed, but Johnson (2004) effectively argued against the evidence for and necessity of systematicity in language and thought, so we do not discuss this feature further.

will vary in many respects based on the details of the individuals and differences in neural states when information makes contact with the nervous system. However, if the individuals tend to share any ecologically specified properties (i.e., if they really are a type) then there will be correspondingly stable aspects of shared neural activity, invariants over the transformation of experience[7]. This subset of neural activity would acquire a certain degree of stability, such that the activity can be re-instantiated, given the right precipitating event, in the absence of the corresponding information. Because ecological information represents properties and not individuals, this kernel of stable neural activity can represent properties associated with a type. This means that ecological neural representations of a certain kind can function as concepts.

This same subset of neural activity can also support *componential* systems. We predict that stable neural representations will only emerge if the corresponding information is sufficiently simple, short, and stereotyped. This type of neural representation is a component – it is a bit that can participate in a number of events made up of other bits. These ecological neural representational components then enable *productivity* in the following way. Neural representations can impact action selection. Some of these actions can be the instantiation of further neural representations. Some of the variance in what actions are selected by a given neural representation will be explained by the learning history of the organism. For instance, if the acoustic event for the letter "B" almost always follows the acoustic event of the letter "A" in a person's learning history, then activating a neural representation for "A" will tend to select the neural representation for "B" (such as when you start singing the ABC song in your head). But some of the variance will also be explained by the current context, summarized in the informational environment and current neural and bodily state of the organism. So, if you are watching West Side Story, then activating a neural representation for "A" may select the subsequent representations "…Jet is a Jet is a Jet all the way." Therefore, ecological neural representations can be combined in multiple ways with other representations and the

[7] We mean this as an explicit analogy to EJ Gibson's (1969) differentiation theory of perceptual learning.

grounded way they do so is functionally related to learning history and current context[8]. Ecological representations can also support componential and productive systems on a deeper level. Because information specifies properties, not individuals, ecological representations don't suffer from the holism that, some argue (e.g. Barsalou, 1999) makes typical perceptual theories unable to support componentiality.

Conceptual, componential, and productive systems support aspects of higher order cognition like counterfactual thinking, thinking about impossible things, and talking about imaginary things. As we said before, we believe that the important task for us here is to show how ecological psychology can support higher-order cognition, whether or not the ecological solution also requires a conceptual and componential system. But, if the reader endorses the logic that aspects of higher order cognition naturally follow from concepts, componentiality, and productivity, then we hope to have shown how ecological neural representations possess these features. We think this demonstration does some important work in justifying the viability of an ecological representational approach, but we would like to add one final point to this discussion. We think that approaching higher cognition from an ecological basis leads to a fundamentally different flavor of analysis to the typical cognitive approach, one which places less emphasis on representational system features reflecting the influence of computer science on cognitive science and more emphasis on action selection and control. We attempt a brief example of such an analysis below.

A common problem that seems to demand mental representation is the act of talking about something imaginary. This is a complex problem if you treat language as a system of reference; when I say the word 'unicorn', to what do I refer? If, instead, you treat language as a system for selecting the actions of yourself and others (i.e. if it is a tool; Bickhard, 2009; Everett, 2012) then this problem becomes identical to the problem of using information conventionally to select an action

[8] It's possible the literature on the limits of classical conditioning and chains of associations may be informative here

(e.g. the bear and the sign example). Our experiences of using the word 'unicorn' dictate the kind of tool that it is and the kind of actions that it can select. When asked to describe a unicorn, a speaker might select the actions 'a horse with a horn' or something similar. From the perspective of informational and neural representations engaged in action selection, talking about imaginary things is exactly the same kind of process as talking about real things, talking about impossible scenarios, considering multiple possible outcomes, and imagining how things might have been different.

The analysis above is very brief and we agree that ecological psychologists should tackle these tricky problems head on, preferably accompanied by data. But, we hope to have shown that recognizing the role of informational representations in action selection and its relationship to the relevant neural representations does provide the necessary foot in the door for an ecological analysis of higher order cognition.

Summary

A representation is a thing that can function in a process 'as if' it were something that process needs but has no access to otherwise; it designates that thing to the process. A good representation is unambiguously (causally) related to its target, making its contents grounded in the contents of the target. Ecological information is such a representation, and it solves two major problems with traditional mental representational systems; symbol grounding and system-detectable error. Together, informational representations and the consequent neural representations can address all major motivations for mental representations including implementing intentionality in a physical system, poverty of stimulus, and representation-hungry 'higher-order' cognition.

Is it worth re-conceptualizing ecological information as a representation, though? After all, ecological psychology has always been explicitly anti-representational and a key feature of the current analysis (ecological information solves the poverty of stimulus problem and thus redefines the job description for the brain) has historically been used to motivate non-representational

approaches to cognitive science (e.g. Chemero, 2009; Wilson & Golonka, 2013). We argue, however, that it is entirely worth it, for two reasons.

First, as we show here for the first time, ecological information is capable of addressing the quite real and pressing motivations typically invoked to support the need for representations, but it does so in a way that solves two major problems. Recognizing the role that ecological information can play in addressing these motivations while avoiding the two problems allows us to build on the hard won theoretical successes from all sides.

The second reason is that treating ecological information as a representation will ensure we use it in our theories for what it really is: a thing that is not the world. Gibson made a mistake calling his approach 'direct perception'. The name made a certain amount of sense at the time, but it has come to imply (for supporters and critics alike) that perception is a free ride; that we simply 'see' the world, no cognitive gymnastics required. This isn't true. Ecological information does specify the world and is therefore a good stand-in for the world, but it is not the same as the world. This means that a) we have to learn to use it as a stand-in (development and perceptual learning should feature front and center of all our theories, e.g. Wicklegren & Bingham, 2001) and b) we will only be able to interact with the world in terms of how it has been projected into information. This last point places meaningful constraints on our search for explanations of our behavior, as in coordinated rhythmic movement (Golonka & Wilson, 2012).

This error, of treating the world as simply a given to a cognitive agent, is quite real. In recent years efforts have been made to extend the ecological approach into domains beyond perception and action, especially language and social psychology (Chemero, 2009; Heft, 2007; Kono, 2009; Schmidt, 2007). These efforts have one thing in common; they all typically extend Gibson's notion of affordances to become opportunities for linguistic or social actions that simply account for the behavior of interest. But, affordances, while interesting, are properties of the world, and must be perceived. The critical question (as Gibson himself emphasized) is actually about whether and how

these properties are informationally specified. We have argued recently (Golonka, 2015) that any extension of the ecological approach into so-called higher-order, ‘representation hungry’ cognition therefore requires extending our understanding of the form, content and use of information, which is what we have continued here. This, we contend, is a representational research program worth pursuing.

References

Agyei, S. B., Holth, M., Weel, F. R., & Meer, A. L. (2015). Longitudinal study of perception of structured optic flow and random visual motion in infants using high-density EEG. *Developmental science, 18(3),* 436-451.

Barsalou, L. W. (1999). Perceptions of perceptual symbols. *Behavioral and Brain Sciences, 22(04),* 637-660.

Bechtel, W. (1998). Representations and cognitive explanations: Assessing the dynamicist's challenge in cognitive science. *Cognitive Science, 22(3),* 295-318.

Bickhard, M. H. (2009). The interactivist model. *Synthese, 166(3),* 547-591.

Bingham, G. P. (1988). Task-specific devices and the perceptual bottleneck. *Human Movement Science, 7(2)*, 225-264.

Bingham, G.P. 1995. Dynamics and the problem of visual event recognition. R. Port & T. van Gelder eds. *Mind as Motion: Dynamics, Behavior and Cognition*: 403–448. Cambridge, MA: 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). Hillsdale, N.J., LEA Publishers.

Bingham, G.P. (2004a). A perceptually driven dynamical model of bimanual rhythmic movement (and phase perception). *Ecological Psychology, 16(1),* 45-53.

Bingham, G.P. (2004b). Another timing variable composed of state variables: Phase perception and phase driven oscillators. In H. Hecht & G.J.P. Savelsbergh (Eds.) *Theories of Time-to-Contact*. Boston: MIT Press.

Breen, M., & Clifton Jr, C. (2013). Stress matters revisited: A boundary change experiment. *Quarterly Journal of Experimental Psychology, 66(10),* 1896-1909.

Bringsjord, S. (2015). The symbol grounding problem… remains unsolved. *Journal of Experimental & Theoretical Artificial Intelligence, 27(1)*, 63-72.

Chemero, A. (2009). *Radical embodied cognitive science*. Cambridge, MA: MIT Press.

Clark, A., & Toribio, J. (1994). Doing without representing? *Synthese, 101(3),* 401-431.

Dietrich, E., & Markman, A. (2003). Discrete thoughts: Why cognition must use discrete representations. *Mind & Language, 18(1),* 95-119.

Everett, D. L. (2012). *Language: The cultural tool.* Vintage.

Fodor, J. A. (1980). Methodological solipsism considered as a research strategy in cognitive psychology. *Behavioral and brain sciences, 3(01)*, 63-73.

Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition, 28(1-2)*, 3-71.

Gibson, J. J. (1966). *The senses considered as perceptual systems.* Boston: Houghton Mifflin

Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton Mifflin.

Golonka, S. (2015). Laws and conventions in language-related behaviors*. Ecological Psychology, 27(3)*, 236-250.

Golonka, S., & Wilson, A. D. (2012). Gibson's ecological approach-a model for the benefits of a theory driven psychology. *Avant, 3(2)*, 40-53.

Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena, 42(1)*, 335-346.

Haugeland, J. (1991). Representational genera. In W. Ramsey, S. P. Stich, & D. E. Rumelhart (Eds.), *Developments in connectionist theory. Philosophy and connectionist theory* (pp. 61-89). Hillsdale, NJ, US: Lawrence Erlbaum Associates, Inc.

Kono, T. (2009). Social affordances and the possibility of ecological linguistics. *Integrative Psychological and Behavioral Science, 43*, 356–373.

Magrassi, L., Aromataris, G., Cabrini, A., Annovazzi-Lodi, V., & Moro, A. (2015). Sound representation in higher language areas during language generation. *Proceedings of the National Academy of Sciences, 112(6)*, 1868-1873.

Michaels, C. F., & Carello, C. (1981). *Direct Perception*. Englewood Cliffs, NJ: Prentice-Hall.

Newell, A. (1980). Physical symbol systems. *Cognitive science, 4(2)*, 135-183.

Pattee, H. H. (2012). Laws and constraints, symbols and languages. In *Laws, Language and Life* (pp. 81-89). Springer Netherlands.

Pylyshyn, Z (1989) Computation in cognitive science. In MI Posner (Ed.) *Foundations of Cognitive Science*. MIT Press.

Ramsey, W. M. (2007). *Representation reconsidered*. Cambridge University Press.

Runeson, S., & Frykholm, G. (1983). Kinematic specification of dynamics as an informational basis for person-and-action perception: expectation, gender recognition, and deceptive intention. *Journal of Experimental Psychology: General*, 112(4), 585.

Schmidt, R. C. (2007). Scaffolds for social meaning. *Ecological Psychology*, 19 (2), 137-151.

Searle, J. (1980). Minds, Brains, and Programs. *Behavioral and Brain Sciences* 3, 417-424.

Steels, L. (2008). The symbol grounding problem has been solved. So what's next? *Symbols and embodiment: Debates on meaning and cognition*, 223-244.

Turvey, M. T., Shaw, R. E., Reed, E. S., & Mace, W. M. (1981). Ecological laws of perceiving and acting: In reply to Fodor and Pylyshyn (1981). *Cognition, 9(3)*, 237-304.

Wallis G., Chatziastros A., Tresilian J. & Tomasevic N. (2007). The role of visual and nonvisual feedback in a vehicle steering task. *Journal of Experimental Psychology: Human Perception and Performance*, 33 (5) 1127-1144.

Wilson, A. D. & Bingham, G.P. (2005a). Perceptual coupling in rhythmic movement coordination - stable perception leads to stable action. *Experimental Brain Research, 164*, 517-528.

Wilson, A. D. & Bingham, G.P. (2005b). Human movement coordination implicates relative direction as the information for relative phase. *Experimental Brain Research, 165*, 351-361.

Wilson, A. D. & Bingham, G.P. (2008). Identifying the information for the visual perception of relative phase. *Perception & Psychophysics, 70(3)*, 465-476.

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.

Wilson, A. D., Snapp-Childs, W., Coats, R., & Bingham, G. P. (2010). Learning a coordinated rhythmic movement with task-appropriate coordination feedback. *Experimental Brain Research, 205(4)*, 513-520.

Van Der Meer, A. L., Svantesson, M., & Van Der Weel, F. R. (2012). Longitudinal study of looming in infants with high-density EEG. *Developmental Neuroscience, 34(6)*, 488-501.

Van der Weel, F. R., & Van der Meer, A. L. (2009). Seeing it coming: infants' brain responses to looming danger. *Naturwissenschaften, 96(12)*, 1385-1391.