bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

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

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Abstract

Cognitive science has three main motivations for assuming that cognition requires representation. These are the need for intentional (meaningful) access to the world, poverty of perceptual access to that world, and the need to support 'higher-order' cognition (e.g. thinking about things in their absence). All representational systems must also address two major problems (symbol grounding and the need for system-detectable error). Mental representations attempt to address the three motivations but stumble over the two problems. 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.

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

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Since the beginning of 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 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. Two major and as yet unsolved problems are the issues of *symbol grounding* (Harnad, 1990) and system-detectable error (Bickhard, 2009). These raise serious objections to the ability of mental representations to support our behaviour, even in principle. In addition, various kinds of 'radical' embodied cognition reject the idea that mental representations are even required (e.g. Chemero, 2009; Gibson, 1966, 1979; Turvey et al, 1981; Michaels & Carello, 1981). Most of these are based in Gibson's (1979) ecological approach to perception and action and use the richness Gibson identified in perception to side-step the need for representations of any kind. However, the major successes of these non-representational accounts have been mostly restricted to domains such as perception-action[1] and they have yet to develop any widely accepted explanations for the 'high-order' cognitive activities driving Motivation 3. As such, they are often ruled out as having the potential to offer a complete account of cognition (e.g. Goldinger, Papesh, Barnhout & Hout, in press).

We believe there might be a way to reframe this debate in a way that preserves the good work of all sides and provides a path for future research. We will agree that cognition requires representations. However, we propose that Gibsonian ecological information perfectly fits the

[1] There have been some non-representational forays into 'higher-order' cognition, e.g. Stephen, Dixon & Isenhower (2009) but these remain few and far between.

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

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.

Why Cognition Must Involve Representations: Three Motivations, a Solution, Two Problems

Newell (1980) proposed a list of constraints any theory of cognition should eventually address[2]. He emphasized the first one in particular, universality, which is the idea that a cognitive system must be able to behave as an (almost) arbitrary function of the environment. In other words, a cognitive system has to be able to be 'about' anything it encounters in the world. He also highlighted that implementing such a flexible, adaptive system in an actual, physical system was a major problem (Newell, 1980; Bechtel, 1998; Stich, 1992). It was argued that universality could only be implemented in a computational system, and that computational systems are necessarily representational (Fodor, 1980; Pylyshyn, 1989)[3]. The need for intentionality therefore provided the first and primary motivation for treating cognition as necessarily representational.

There are two other commonly cited motivations in which representations are required to fill apparent gaps in the causal chain of events leading from world to functional behavior in that world. Motivation 2 is that representations are required to overcome a gap from world to sensory receptors, caused by a poverty of stimulus (e.g., Haugeland, 1991). This motivation reflects the widespread assumption that sensory input "represent[s] but few, and biologically unimportant, characteristics of objects" (Gregory, 1968 p 279). Motivation 3 is that representations are required

[2] From his Table 1, the constraints are, in order: 1) Behave as an (almost) arbitrary function of the environment (universality); 2) operate in real time; 3) exhibit rational, i.e. effective adaptive behaviour; 4) use vast amounts of knowledge about the environment; 5) behave robustly in the face of error, the unexpected, and the unknown; 6) use symbols (and abstractions); use (natural) language; 8) exhibit self-awareness and a sense of self; 9) learn from its environment; 10) acquire its capabilities through development; 11) arise through evolution; 12) be realisable within the brain as a physical system; 13) be realisable as a physical system.
[3] But see Picninini 2004, 2008 for an argument against the necessity of representation in computation

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

to cross a second gap, from sensory receptors to functional behavior. This motivation is about supporting 'higher-order' cognition that isn't obviously driven by immediate sensory states (so called 'representation-hungry' problems such as thinking about things in their absence; Clark & Toribio, 1993). We take Motivation 1 (getting intentionality out of a physical system) to be the primary job of representations. Motivations 2 and 3 are constraints on exactly how Motivation 1 might be implemented given the existence of the two gaps.

The current kind of representations that cognitive science uses to addresses these three motivations are *mental representations*. These are typically considered to be computational systems of some kind or another because such systems have the necessary properties to allow intentionality, enrich impoverished stimuli via inference and support higher order cognition (Newell, 1980; Pylyshyn, 1989; Fodor, 1980). There are certain concerns with mental representations that remain unsolved, however. The first is the *symbol grounding problem* (Bickhard, 2009; Harnad, 1990; Searle, 1980); the proposal that symbolic or representational systems do not necessarily come with intrinsic access to intentional content and thus cannot actually be 'about' anything without additional help[4]. The second, related problem is the *system-detectable error problem* (Bickhard, 2009), which is that ungrounded 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. These problems are not as widely discussed in the literature as they were in the 1990s, but they remain as-yet unsolved issues for mental representations and these therefore remain a problematic solution to modelling cognition.

We now have a job description (for cognitive representations; they must address all three motivations and not fall prey to the two problems. The remainder of this paper will lay out our argument that ecological representations (informational representations and the neural

[4] Bickhard (2009) frames this as the need for representational content to be *internally related* to the thing it represents. This means that the relationship between a representation and its contents must be defined within the bounds of the representation, and cannot rely on an external source.

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

representations they create) fill this job description extremely well, up to and including the daunting question of supporting 'higher order' cognition. This last section is a continuation of the argument begun in Golonka (2015) about how to extend the ecological approach beyond its traditional territory of perception and action.

Representations: A Definition

While there is considerable disagreement about many of the details of representations, most cognitive scientists (pro or con) endorse the idea that a representation is a thing that *stands in for* something else[5]. Newell (1980) called this property *designation* and defined it 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).

Some process P requires access to Y but there is a gap between them. A representation, X, is a thing that is not Y but can close the gap and that P can access and use as if it were Y; when it does, P works as if it had access to Y. The net result is 'action at a distance', where the distance comes from whatever caused the gap between P and Y. Taking this as a minimal definition of representation has the advantage of being neutral with respect to the separate question of representational format,

[5] Bickhard (2009) argues that any representational system based on stand-ins (encodings) alone is insufficient to explain cognition because these systems lack a mechanism for representations to emerge and for the system to detect errors. We address how we solve these two problems below. However, Bechtel (1998) notes that there is widespread agreement that representations should have the function of standing in; that is, natural signs (Hatfield, 1990) and indexes (Dretske, 1988) wouldn't count as representations since their ability to stand in for a state of affairs is epiphenomenal. Representations as stand-ins also remains the dominant paradigm in cognitive psychology, so this is our focus.

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

e.g., whether or not they are compositional (Fodor & Pylyshyn, 1988), dynamic (Freyd, 1987), propositional (Pylyshyn, 1981) or depictive (Barsalou, 1999; Kosslyn, 1995).

We should mention that Newell's is a fairly minimal definition of representation and there are authors who add additional constraints when defining representations used by cognitive systems. Haugeland (1991) includes poverty of stimulus as a constituent of the definition of representations, rather than simply a pre-condition that makes them necessary, while Dietrich and Markman (2003) require that cognitive representations must be internal to cognitive systems. The argument we present below, about ecological information as representation, directly challenges the assumption of poverty of stimulus and concerns a step prior to anything getting into the organism. Therefore, our account kicks in before such additional concerns become relevant, and so we will present our argument with respect to Newell's basic definition and then examine the consequences, specifically whether ecological representations can fill the job description for a cognitive representation.

Ecological Information is a Representation

Ecological information takes the form of higher order relational patterns in energy arrays (e.g. the optic array). These patterns are created by the lawful interaction of the energy with the dynamics of the world (Turvey, Shaw, Reed & Mace, 1981) and are used by organisms to perceive that world (Chemero, 2009; Gibson, 1979). In order to identify whether information is a representation, we need some details about the process that creates it.

Properties of objects and events in the world change over both space and time in particular ways that reflect the composition and organization of these properties. The appropriate level of description for adequately capturing this composition and organization is the level of *dynamics* (Bingham, 1995). Formally, dynamics is a level of description of how things change over time that includes reference to the forces involved; the units are those of time, position (and its derivatives) and mass.

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Successfully interacting with a task means organizing behavior so as to complement these task dynamics. Organisms do not have direct access to task dynamics. Most of the behaviorally relevant dynamics in the world are 'over there' and not in mechanical contact with the organism. They must therefore be perceived (Gibson, 1979). Perception relies on information about dynamics, but information (specifically, Gibsonian ecological information) is only a *kinematic* projection of those dynamics into an energy array (Bingham, 1988; Turvey et al, 1981). Formally, kinematics is a level of description that only refers to motions; the units are time, position (and its derivatives) but not mass. This means that kinematic information cannot be identical to the dynamical world, and this fact is effectively a poverty of stimulus (Bingham, 1988 refers to it as the 'perceptual bottleneck'). Despite the fact that the information and the world cannot be identical, the lawful process that projects the dynamical world into kinematic, ecological information does result in information variables that *specify* (i.e., map 1:1 to) the dynamics that caused them (Gibson, 1979; Runeson & Frykholm, 1983; Turvey et al, 1981).

A simple example of this idea comes from research on coordinated rhythmic movement. In this task, people try to move their limbs so as to preserve a specific mean relative phase between them; for example, moving two fingers up and down so that they are doing the same thing at the same time (0° mean relative phase) or the opposite thing at the same time (180°). Controlling a behavior entails perception, and the relevant dynamical 'world' property to be perceived in this task is *relative phase.* Recent ecological research on this task has identified that the kinematic perceptual information that specifies relative phase is *the relative direction of motion* in the local optic flow caused by the moving limbs (Wilson & Bingham, 2008). People move so as to make relative direction behave in a certain way, and, because it specifies relative phase, each behavior of relative direction produces one kind of coordinated rhythmic movement. Relative direction (information) is not the same as relative phase (world property) however, and this has consequences. For example, there is no particular difference between 0°, 90° and 180° at the level of relative phase; but at the level of relative direction, 0° is all common motion while 180° is all relative motion, and the optical motion at

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

90° is maximally variable (half the time common, half the time relative, and constantly switching). This is the reason why behaviourally, 0° is more stable than 180° while 90° is maximally unstable; we interact with the dynamical world via kinematic information (e.g. Wilson, Collins & Bingham, 2005a).

Something immediately leaps out from this analysis. The kinematic information variable 'relative direction' is *standing-in* for the dynamical world property 'relative phase' and it requires no additional enrichment from a mental representation in order to do so. Consequently, in line with the designation definition above, we propose that ecological information simply *is* the representation that closes the poverty of stimulus gap, though it is *external* and *ecological* rather than *internal* and *mental*.

Let us unpack this idea with reference to the definition of designation given above. Ecological information is a thing, X (a time-extended kinematic pattern in an energy medium, e.g. relative direction) that is not another thing, Y (a dynamical property of a biologically relevant object or event in the environment, e.g. relative phase) but that designates Y to a process P (an organism that needs to organize a functional response to the environmental properties, e.g. a coordinated rhythmic movement). For perceiving-acting organisms, the function of information is to stand-in for the world. That is, the consequence of ecological information on nervous systems is best explained by invoking the ability of information to represent properties of objects and events in the world, and it happens to be a good (effective) representation because of specification.

Why does ecological information get to claim a 1:1 mapping between the world and an energy distribution while cognitive sensation-based accounts do not? Why is it not merely an interesting but fundamentally ambiguous cue? The key difference is that the cognitive analysis is *extensional*, while the ecological analysis is *intensional* (an argument first detailed by Turvey et al, 1981; see also Gibson & Gibson, 1955, for an early discussion of the difference). Briefly, the extensional cognitive analysis of what properties are projected into an energy array begins with sets of *objects* that have a certain property (e.g., the set of sit-on-able objects). It can then be shown that some of the items in

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

this set are co-extensive with items in another set of objects that have some other property (the set of lie-on-able objects). In other words, there is a set of objects that possess both properties. Because the objects are co-extensive, so too are the sensory consequences of those objects. Therefore, from this extensional starting point, you cannot perceptually specify (map 1:1) properties of objects and you need some additional, internal process to resolve the ambiguity. The intensional ecological account, in contrast, begins by defining a biologically relevant *property* of an object or event (e.g., the affordance that something can be sat on) and then identifying a pattern in an energy array that is unique to that property. Rather than inferring properties on the basis of object identification, the ecological approach suggests that perception is about perceiving the properties themselves. This means that even if multiple objects share the same property, if the property creates information then it is still perceptually specified (mapped 1:1) every time it occurs.

This particular 'non-identical-yet-1:1' mapping between objects and events in the world and ecological information fits Newell's definition of a representation, and it's a good one to boot. It therefore seems perfectly justifiable to refer to an ecological information variable as an ecological representation of some property of the world, a representation that fills the causal gap between behaviorally-relevant but distant properties of objects and events and the nervous systems of perceiving/acting organisms.

Informational Representations Create Neural Representations

So far we've discussed informational representations in terms of their role as external representations of behaviorally-relevant properties of the environment. This crosses the first gap (from world to perceptual systems). This leaves us with the second gap (from perceptual systems to functional behavior), which is where mental representations typically show up in standard cognitive theories. Perhaps the neural activity caused by interacting with an informational representation deserves to be called a representation as well, and perhaps this neural representation will still be like the mental representations of standard cognitive science? This is an alluring possibility

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

considering that it sounds an awful lot like these internal representations would essentially be perceptual symbols, in the sense developed by Barsalou (1999). We think that this possibility disappears upon closer inspection and we address this after we define our notion of a neural representation.

First, let's see how neural activity caused by ecological information stacks up to Newell's definition of a representation. According to this definition, "having X (the symbol) is tantamount to having Y (the thing designated) for the purposes of process P" (Newell, 1980, p. 156). Neural activity precipitated by ecological information would be an internal representation if such activity X, which is not another thing, Y (i.e. that information variable), designates Y to a process P (allowing behavior to be coordinated with respect to the information specifying the relevant property in the world) and allows P to exhibit action with respect to Y at a distance (the distance between perceptual systems and the rest of the body). If these conditions are met, it would be fair to say that external informational representations cause internal neural representations.

This seems to us to be perfectly possible, but whether this is what happens is actually an unanswered (although empirically testable) question. Explicitly ecological analyses of neural activity (i.e. tracking the neural consequences of interacting with ecological information) are currently rare at best, although work by Agyei et al (2015), van der Meer et al (2012) and van der Weel & van der Meer (2009) are the kind of programs we have in mind. Other lines of neuroscientific enquiry do suggest that at least some of the structure of energy impinging on perceptual receptors is preserved as it travels through the nervous system (e.g. Magrassi et al, 2015). There is also a rich collection of behavioral work from the ecological approach showing that the form of a behavior maps tightly onto the structure of the information involved, suggesting that structure is carried through the nervous system (e.g. the fact that relative phase is specified by relative direction accounts for most of the behavioral level characteristics of coordinated rhythmic movement; see Bingham, 2001, 2004a, b; Golonka & Wilson, 2012; Snapp-Childs, Wilson & Bingham, 2015; Wilson, Collins & Bingham, 2005a,

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

b; Wilson, Snapp-Childs & Bingham, 2010). 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, and since the structure of ecological information is objectively verifiable, it can guide future work investigating the extent to which this structure is propagated or systematically transformed by neural activity.

So, do these neural representations put us back on familiar ground in cognitive science? We think not. Any psychological power possessed by these neural representations comes from their resemblance 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.

Why Replace Mental Representations with Ecological Representations?

We have identified three non-trivial reasons motivating the hypothesis that cognition requires representations. We need to get intentionality out of a physical system; we have to cross the gap from world to sensory systems; and we need to cross the gap to functional behavior in the absence of currently available information. These concerns led to the hypothesis of computational mental representations as the kind of representation cognition has available to address these motivations. However, any representational system must contend with two critical issues (symbol grounding and the need for system detectable error) and there is a case to be made that no current cognitive representational framework adequately does so (Bickhard, 2009). In the following sections, we will detail how ecological representations (informational and neural) solve both problems and also

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

address all three motivations, making them better candidates for the representations needed for cognition.

Ecological Information Solves Two Problems with Mental Representations

Ecological Information Solves the Symbol Grounding Problem: The symbol grounding problem[6] facing mental representations is familiar to most and has been discussed at length (e.g., Harnard, 1990) so we will characterize the problem only very briefly. Symbolic representations suffer from their arbitrary mapping between the structure of the representation and the thing it's meant to represent. If functional behavior arises through the successful manipulation of arbitrary symbols on the basis of syntax alone, then meaning is entirely extrinsic to the system. In Bickhard's terms, the representation is externally defined, which means that the system effectively only "knows" how to implement rules that operate over symbols and not what those symbols mean (typically demonstrated with the Chinese Room argument; Searle, 1980).

According to Harnard (1990), 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) because higher level systems inherit the lower level grounding. Proponents of standard cognitive representational approaches are therefore happy that structures that are lawfully related to objects and events in the world do not suffer from a symbol grounding problem. Attempts to solve the symbol grounding problem therefore try to link abstract or higher-order mental representations to more basic perceptual representations (e.g. Barsalou, 1999; Harnard, 1990). However, this does not solve the problem for standard cognitive psychology, because it treats perception extensionally, which means even the lower perceptual representations lack internally defined content and there is no grounding for the higher processes to inherit. Bickhard (2009) suggests that this "infinite regress of interpreters interpreting" (p 573) is endemic to any

[6] Not everyone thinks the problem remains; see Steels, 1997, and a more recent solution from Taddeo & Floridi (2007). But the debate remains (e.g. Bringsjord, 2015) and so, for the people who, like us, still worry about this problem, we endeavour to solve it.

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

representational account where content is defined externally. Bickhard argues that familiar mental representational accounts (he includes arguments against Milikan, Fodor, Dretske, and Cummins) all define representations externally. Therefore, none of these mental representations are grounded and their contents are remain inaccessible to the organism.

Informational representations, however, are immune to the grounding problem. This is because 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: The second problem, system detectable error, follows from the first (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.

First, ecological information (the representational vehicle) is internally related to events in the world (the representational content); it is grounded. Second, ecological information can be used to coordinate and control the production of stable, functional behavior. If the information being used cannot support such behavior (because it is the wrong variable to use, or because of an experimental perturbation) behavioral control will fail; an error will be incurred. Because, from an ecological perspective, perceiving and acting are fundamentally intensional (Gibson, 1979, Turvey et al, 1981) and because the content of informational representations is accessible to the perceiving-

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

acting organism, organisms can be aware of when these representations are wrong and this awareness can have a consequence on future behavior.

For example, as noted above, the relative phase of a coordinated rhythmic movement is typically perceived using relative direction as the information (Bingham, 2001, 2004a, b; Wilson & Bingham, 2008; Wilson, Collins & Bingham, 2005b). Relative direction is stable and clearly detected at 0° and 180° and so these coordination behaviors are stable as well (e.g. Wilson, Collins & Bingham, 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). If that variable is then perturbed (i.e. experimentally made uninformative) then the trained performance at 90° becomes unstable again (Wilson & Bingham, 2008). In our new, representational terminology, relative direction is a representation of relative phase that enables stable behavior at 0° and 180° but not 90°. This instability is detected and drives perceptual learning of a new representation, relative position, which now enables stable behavior at 90°. The perturbation disrupts the intensional nature of relative position; it exists and is detected but it no longer points to relative phase, and behavior fails again.

Ecological Information Addresses the Major Motivations for Representations

Information fits the job description of a representation, and solves two important problems faced by any representational system. Here we describe how ecological information addresses the three inter-related motivations for why cognition must be representational. Addressing the first two motivations falls out of the nature of ecological information. Specifying information is intrinsically intensional (Turvey et al, 1981) and is, of course, far from impoverished (Gibson, 1966, 1979).

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Addressing the third motivation about 'higher order cognition' will require some new analysis (begun in Golonka, 2015 and continued here).

Motivation 1 is that representations are required to connect the 'knowledge-level' to the 'physical-level' descriptions (Newell, 1980; Pylyshyn, 1989). Because informational representations specify biologically relevant task dynamical properties (rather than perceptual primitives), they are inherently meaningful (Turvey et al, 1981). Thus, intentional behavior follows from coordination with informational representations. Second, because information specifies properties and not individuals (Turvey et al, 1981) informational representations can explain our behavioral flexibility. When we encounter a novel object or event, it will likely project at least some familiar information variables (e.g., whether it is moveable, alive, etc), giving us a basis for functional action in a novel context. Because information also solves the symbol grounding and system-detectable error problems (see below), it also supports error-driven learning to refine that initial response into something task appropriate.

Motivation 2 is that representations are required to bridge a poverty of stimulus gap. Ecological information crosses the gap between properties of the world and our sensory receptors because of the kinematic specification of dynamics (Bingham, 1995; Gibson, 1979; Runeson & Frykholm, 1983; Turvey et al, 1981). Kinematic information is not identical to (and slightly impoverished relative to) the dynamical world but it can specify it and the lawful process that underpins the projection of dynamics into specifying kinematics allows information to cross the gap in an effective and reliable manner. While there is a physical distance between us and much 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.

We now turn to the main event; *Motivation 3 is that representations are required to explain behavior shaped by things that are not perceptually specified right now.* This is, to put it mildly, the motivation ecological accounts have yet to tackle in detail. Most (if not all) research on ecological

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

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). We can think of things in their absence. We can think of things that could, but might not, happen. We can think of impossible things. Furthermore, these thoughts appear to influence subsequent behavior. One might object to our analysis so far on the basis that it hasn't addressed these, more conceptual, aspects of our behavior (e.g. Goldinger, et al, in press for an unambiguous example of this in action).

Even Barsalou (1999), who worked to ground representations in perceptual experiences, hypothesized that those experiences are later reified into perceptual symbols which could form the basis of "higher" cognitive functioning. Barsalou and others all assume that knowledge systems must be both *conceptual* (cf. Haugeland, 1991) and *componential* to allow complex expressions to be decomposed and new expressions to be built up (e.g., Dietrich & Markman, 2003; Fodor & Pylyshyn, 1988); that is, to enable *productivity* (Chomsky, 1957; Fodor & Pylyshyn, 1988). 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[7]. 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

[7] One other feature that gets discussed is *systematicity*, 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.

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

system. However, as it turns out, we think that certain features of the two ecological representations mean that they enable conceptual, componential systems.

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 (as in in the production of coordinated rhythmic movements, e.g. Bingham, 2001, 2004a, b). 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 any particular 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 the 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

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

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[8].

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 is getting experienced drivers to mime steering. Despite the fact drivers can successfully steer a real car, they are unable to realistically mime the action of steering, and often do so in ways that would have catastrophic consequences in actual driving contexts (e.g. Wallis, Chatziastros, Tresilian & Tomasevic, 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.

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

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

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

"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."[9] 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. 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.

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

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

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 Newell's definition of a representation with respect to corresponding ecological information will be considered a neural representation. Second, let us recall that not all neural 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.

First, let us tackle how such neural representations might function as components in conceptual, componential, and productive systems. It is uncontroversial to say that developing a concept requires experience with multiple individuals of a type. From our perspective, this development 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 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 neural activity. 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

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

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 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 grounded way they do so is functionally related to learning history and current context. 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

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

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 (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

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

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. 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 reconceptualising 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

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

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) 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, e.g. understanding how relative phase is perceived was crucial to understanding and modelling the behavioural characteristics of coordinated rhythmic movement (see 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.

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

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. *The 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. (2003). An outline of a theory of affordances. *Ecological Psychology, 15* (2), 181-195.

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

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

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

Craver, C. F. (2007). *Explaining the brain*. Oxford University Press.

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

Dretske, F. I. (1988). *Explaining behavior: Reasons in a world of causes*. Cambridge, MA: MIT press.

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.

Gibson, J. J., & Gibson, E. J. (1955). Perceptual learning: Differentiation or enrichment? *Psychological review*, *62*(1), 32.

Goldinger, S. D., Papesh, M. H., Barnhart, A. S., Hansen, W. A., & Hout, M. C. (in press). The poverty of embodied cognition. *Psychonomic Bulletin & Review.*

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).*

Gregory, R. L. (1968). Perceptual illusions and brain models. *Proceedings of the Royal Society of London. Series B, Biological Sciences, 171*(1024), 279-296.

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

Hatfield, G. (1990). Gibsonian representations and connectionist symbol processing: Prospects for unification. *Psychological Research, 52*(2-3), 243-252.

Haugeland, J. (1991). Representational Genera. *Philosophy and connectionist theory, 61*.

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

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Kosslyn, S. M. (1981). The medium and the message in mental imagery: A theory. *Psychological review, 88*(1), 46.

Kosslyn, S. M. (1995). Mental imagery. *Visual cognition: An invitation to cognitive science, 2*, 267-296.

Lees, R. B., & Chomsky, N. (1957). Syntactic Structures. Language, 33(3 Part 1), 375-408.

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* (pp. 1-208). 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. W. (1981). The imagery debate: Analogue media versus tacit knowledge. *Psychological review, 88(*1), 16.

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

Silberstein, M., & Chemero, A. (2013). Constraints on localization and decomposition as explanatory strategies in the biological sciences. *Philosophy of Science*, *80*(5), 958-970.

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

Stephen, D. G., Dixon, J. A., & Isenhower, R. W. (2009). Dynamics of representational change: Entropy, action, and cognition. *Journal of Experimental Psychology: Human Perception and Performance*, *35*(6), 1811.

Taddeo, M. & Floridi, L. (2007). A Praxical Solution of the Symbol Grounding Problem. *Minds and Machines 17(4)*, 369–389.

bioRxiv preprint doi: https://doi.org/10.1101/058925; this version posted June 15, 2016. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

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. (2005a). 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.