A Potential Mechanism for Gibsonian Resonance: Behavioral Entrainment Emerges from Local Homeostasis in an Unsupervised Reservoir Network

J. Benjamin Falandays[1], Jeffrey Yoshimi[2], William Warren[1], & Michael Spivey[2]
[1]Brown University, Dept. of Cognitive, Linguistic, and Psychological Sciences
[2]University of California, Merced, Dept. of Cognitive and Information Sciences

Abstract

While the cognitivist school of thought holds that the mind is analogous to a computer, performing logical operations over internal representations, the tradition of ecological psychology contends that organisms can directly "resonate" to information for action and perception without the need for a representational intermediary. The concept of resonance has played an important role in ecological psychology, but it remains a metaphor. Supplying a mechanistic account of resonance requires a non-representational account of central nervous system (CNS) dynamics. We present a series of simple models in which a reservoir network with homeostatic nodes is used to control a simple agent embedded in an environment. This network spontaneously produces behaviors that are adaptive in each context, including (1) visually tracking a moving object, (2) substantially above-chance performance in the arcade game *Pong*, (2) and avoiding walls while controlling a mobile agent. Upon analyzing the dynamics of the networks, we find that behavioral stability can be maintained *without* the formation of stable or recurring patterns of network activity that could be identified as neural representations. These results may represent a useful step towards a mechanistic grounding of resonance and a view of the CNS that is compatible with ecological psychology.

Keywords: ecological psychology, resonance, action-perception, computational cognitive neuroscience, reservoir computing

A Potential Mechanism for Gibsonian Resonance: Behavioral Entrainment Emerges from Local Homeostasis in an Unsupervised Reservoir Network

Introduction

Thinking of the mind as analogous to a computer was a key inspiration for many thinkers central to the founding of cognitive science as a field some 50 years ago, and remains a popular notion today. For these cognitivist thinkers, cognition is a process of performing logical operations over internal representations that stand for entities and ideas. This view of cognition can be traced back at least to the psychophysics work of Hermann von Helmholtz in the mid 19th century (1860), who first popularized the notion of perception as inference. Cognitive agents, Helmholtz thought, have direct access only to their own sense data, but not to the things in the world that cause sense data, and therefore must infer the latter from the former. In philosophy of mind, this approach has been referred to as indirect or representational realism.

But for as long as this stance has dominated conceptions of mind and brain, it has also had its detractors. Many have argued that cognitivism introduces a false dualism between stimulus and response, and mistakenly paints the organism as a passive entity (Dewey, 1896). Consider that in 1942, before the advent of modern computing technology, a different metaphor was commonly used to express a cognitivist stance: the brain was said to be like a (musical) keyboard, on which external stimuli would play (through sensory impulses) to produce melodies “depending on the order and the cadence of the impulses received” (i.e. neural and subsequent behavioral responses). The phenomenologist Merleau-Ponty (1942) took issue with the keyboard metaphor, writing:

> “The organism cannot properly be compared to a keyboard on which the external stimuli would play [...] for the simple reason that the organism contributes to the constitution of that form [...] When the eye and the ear follow an animal in flight, it is impossible to say "which started first" in the

> exchange of stimuli and responses [...] since all the stimulations which the organism receives have in turn been [made] possible only by its preceding movements which have culminated in exposing the receptor organ to the external influences [... I]t is the organism itself—according to the proper nature of its receptors, the thresholds of its nerve centers and the movements of the organs—which chooses the stimuli in the physical world to which it will be sensitive [...] *This would be a keyboard which moves itself in such a way as to offer—and according to variable rhythms—such or such of its keys to the in-itself monotonous action of an external hammer*[1]."

In this passage, Merleau-Ponty attempted to revise the standard metaphor of his time, presenting cognition not as passive process driven by the environment, but instead as an active one, driven mutually by organism and the environment—akin to a keyboard that is both played and plays itself by pressing its keys onto the world around it.

A similar line of argumentation was prominently taken up by the school of thought known as ecological psychology, associated with James and Eleanor Gibson, and more recently in the framework of Embodied Cognition. Researchers in these traditions argue that the cognitivist approach introduces an insurmountable chasm between mind and world, making it impossible for cognitive agents to ever access the meanings or referents of their internal representations (Michaels & Palatinus, 2014). Gibson emphasized that perceiving-acting organisms have no need to represent the world outside, and instead can "resonate" to structured flows of energy—an idea he called "direct perception." For example, a bee attempting to fly through a small gap need not build up an internal representation of the environment, its own body, and calculate a trajectory. Instead, it could solve the problem simply by moving in such a way that the speed of image movement in the left and right hemifields is balanced in the right and left eyes, which will ensure the bee passes through the center of the gap (Srinivasan, 1992; Duchon & Warren Jr, 2002

---

[1] Emphasis ours

found evidence that humans use the same strategy). Thus, rather than compute, the bee can “resonate” or “attune itself” to information that uniquely specifies useful relationships between action and perception, known as "affordances" (Chemero, 2003).

While the tradition of ecological psychology has produced many important insights, in eschewing the notion of representation and instead focusing on what goes on at the level of the organism and environment, this tradition has avoided the issue of how neural activity figures into the story. This was an important move in order to call attention to the fact that many cognitive problems need not require complex internal representations. However, recently there have been increasing calls to finally reintroduce neural dynamics into ecological theories of cognition, towards fleshing out a mechanism for Gibson’s notion of resonance, which remains a metaphor (Raja, 2018, 2019, 2021; de Wit & Withagen, 2019).

We suggest that a useful step towards such a mechanistic account of resonance is to emphasize the role of homeostatic neural mechanisms in facilitating self-organization of the CNS. Ecological psychology has focused on homeostatic mechanisms (often described as "control laws") at the level of the organism-environment relation, while work stemming from the cybernetic tradition has emphasized how internal homeostatic mechanisms can lead to adaptive behavior at the system level (Ashby, 1960). It is well established that the CNS implements several different homeostatic mechanisms, including synaptic scaling and regulating the expression of ion channels (Desai, 2003; Chistiakova, Bannon, Bazhenov, & Volgushev, 2014; Turrigiano & Nelson, 2004; O’Leary & Wyllie, 2011), which allow the CNS to stabilize activity following perturbations (see also: Grossberg, 1982). We propose that the emergence of adaptive behavior at the organism-environment level by virtue of homeostatic mechanisms in the CNS constitutes a viable mechanistic account of Gibson’s concept of resonance.

In this paper, we present a series of simple models that serve as a proof-of-concept that homeostatic properties of the CNS can generate adaptive behavior at the organism-environment level[2]. The central component of our models is an sparsely-connected, spiking neural network—a reservoir computer—composed of homeostatic nodes. These nodes adjust connection weights with neighbors and "target" activity levels in order to keep their flow of activity consistent over time. This model was first introduced by Falandays, Nguyen, and Spivey (2021), where it was applied to the context of language processing. When the network was fed inputs generated by a simple probabilistic grammar, it was shown to exhibit behavioral signatures associated with "predictive coding," including increased activity for surprising inputs and sequence completion, suggesting that the network is able to entrain itself to complex sensory patterns that unfold over time, without the need for supervision. We hypothesized that the same network would serve to control behavior when embedded in an action-perception loop, while avoiding some of the representational assumptions of cognitivism.

With short periods of unsupervised training, we observe that the model produces adaptive behavior in a variety of contexts, including spontaneous object-tracking behavior (following a moving stimulus despite no explicit instruction to do so), above-chance performance in the arcade game *Pong*, and wall-avoidance behavior. The reservoir activity underlying these behaviors can be seen as a simple illustration of Gibson's notion of resonance, and offers an opportunity to consider how internal (neural) dynamics and movement work together in this phenomenon. We explain this as multi-scale resonance, whereby individual nodes resonate to flows of energy in their immediate environment, which in turn drives movement and new perceptions, ultimately allowing the agent as a whole to resonate to information in the external environment.

In what follows, we begin with some background on the contrast between

---

[2] We emphasize that our proposal is a form of "how-possibly" explanation (Dray, 1968): our model shows *one possible* mechanism by which resonance could occur, but much further work will be necessary to determine if, or to what extent, something like this mechanism *actually* accounts for the behavior of humans or other organisms.

representational and direct-realist approaches to cognition by virtue of an oft-cited example, the “outfielder problem.” We use this example to clarify key tenets of ecological psychology, including the notions of resonance and attunement. Then, we reflect on the lack of a mechanism for Gibson’s concept of resonance. We explain why such a mechanism is needed, why standard views in cognitive neuroscience are not up to the task, and discuss some proposals for how to fill this gap. Next, we offer our own proposal for a system with the potential to ground the concept of resonance in the CNS—the reservoir network—and suggest that endowing these networks with self-stabilizing mechanisms is an important step forward. Then, we present our model and analyze its behavior in three agent-environment systems, demonstrating that apparently-adaptive behavior at the agent level emerges from the homeostatic mechanism at the level of nodes, and is not dependent upon the formation of stable and/or recurring activity patterns of the kind that might be expected within a representational theory of CNS.

A primer on the ecological approach to perception-action

Newcomers to ecological psychology may find themselves a bit overwhelmed by the prevalence of jargon associated with the field. Given the major differences between this approach and the more dominant cognitivist tradition with which readers may be more familiar, ecological psychologists have found it necessary to introduce a number of new terms. Many of these terms have proven crucial for theory-building in this tradition, so to not deter the uninitiated, this section will provide a brief primer on the direct-realist approach to cognition and define some key concepts.

The outfielder problem is a classic example used to illustrate the differences between representational and direct-realist approaches to perception and action. In the outfielder problem, a baseball player must view a fly ball and decide where to run in order to catch it. A representational approach to this problem would involve the player’s brain constructing a mental representation of the fly ball’s trajectory, based on visual input and other sensory information. This mental representation would then be used to guide the player's actions, such as running towards the predicted landing point of the ball (Saxberg, 1987).

On the other hand, the direct-realist suggests that no such mental representation is needed. Instead, the player could use continuous, online visual information to guide their movements to the right place at the right time. For example, Fink, Foo, and Warren (2009) provided evidence that outfielders control running direction and speed so as to cancel the optical acceleration of the ball, which results in intersecting at the landing location at the right time.

In this example, we can say that the player's actions have become perceptually coupled to a pattern of *optic flow*, meaning a pattern of change in light hitting the retina, due to the relative motion of an individual and objects in the environment. Ecological psychologists use the term "optic flow" in this case, rather than simply "visual information," because the former emphasizes (1) a pattern of sensory stimulation over time, rather than static imagine in a slice of time, and (2) that visual information is generated both by changes in the environment and by the motion of the observer.

Going beyond the context of vision, observer-relative patterns of sensory change have been referred to as "ecological information," to distinguish this notion from other uses of the term "information." A more common use of the term "information" among cognitive scientists is the one used in information theory, which is operationalized as the reduction of uncertainty, or surprise, upon receiving a signal. Ecological psychologists emphasize that this more common notion of information is purely syntactic, meaning it deals only with the relationships among arbitrary signals, stripped of all semantic content.

In contrast, ecological information is inherently *semantic*, in that it specifies the state of the animal-environment system, and thus has meaning or value for an organism. For example, Gibson argued that patterns of optic flow can directly specify opportunities for action—called "affordances"—and that adaptive action involves the perception of these affordances. It is in this sense that Gibson thought perception was *direct*: organisms perceive useful relationships between themselves and the environment by detecting information that uniquely specifies what can be done, instead of needing to make inferences about the environment based on impoverished, static sensory information.

When an organism perceives affordances for action, Gibson described the organism as "resonating" to ecological information (M. Anderson & Chemero, 2019). The concept of resonance draws an analogy between the way that an organism becomes coupled to their environment, and the way that two nearby tuning forks, for example, may become physically coupled by sound waves. Consider that each tuning fork (and indeed any object) has a "resonant frequency," a natural frequency of vibration due to its physical properties (e.g. size, shape, material). When one tuning fork is struck and begins to vibrate, sound waves may travel to a nearby tuning fork, causing the latter to vibrate as well, *if* the two have the same resonant frequency. In this vein, Gibson's use of "resonance" to describe the detection of information implies something like a spontaneous physical coupling of two parts of a system by virtue of energy transferred through a physical medium.

The metaphor of resonance can be expanded upon to describe the ecological view of learning. Consider that the resonance of our two tuning forks requires them to share the same resonant frequency—the natural frequency of oscillation of an object, determined by its physical properties. When resonant frequencies match, we may say that one tuning fork is "attuned" to the resonant frequency of the other. The resonant frequency of an object can be altered through physical changes, such as clipping a damper to a tuning fork, or adjusting the tension on a guitar string. Along these lines, Gibson described learning as analogous to attunement—the altering of parameters (e.g. visual-system parameters) so as to resonate to information for an affordance. For example, a novice ballplayer may not be aware that running so as to cancel optical acceleration of a ball will lead them to the landing point, but through experience they may gradually adjust their perception-action system to detect and cancel optical acceleration, becoming attuned to this information. Gibson described this view of learning as being about "differentiation"—the gradual refinement and calibration of existing perception-action capacities—as opposed to "enrichment," or the *adding* of knowledge or new mental capacities (Gibson & Gibson, 1955).

Finally, ecological psychologists emphasize that resonance enables one individual to *anticipate* the behavior of another individual or object. For example, by resonating to the appropriate information, our ballplayer is able to anticipate the motion of the ball, going towards where it *will* be. The notion of anticipation used here can be distinguished from the term "prediction," where the latter involves a mental model of a target's behavior and the formation of an explicit expectation about what will happen (Falandays et al., 2021; Bickhard, 2016; Zhao & Warren, 2015; Stepp & Turvey, 2010).

What is the mechanism of resonance?

One strength of Gibson's concept of resonance is that it treats cognition (i.e. perception-action) as a kind of physical coupling, implying that we need not invoke intermediate representations or symbolic operations. In the spirit of this idea, ecological psychologists have tended to focus on explanations that lie at the level of organism-environment interactions, down-playing the role of the brain. But, Raja and colleagues have recently drawn attention to the fact that Gibson left the concept of resonance as a metaphor (Raja, 2018, 2019, 2021; de Wit & Withagen, 2019). Humans, of course, are not tuning forks, so what does it actually mean for *us* to resonate to ecological information? Once we commit to the idea that perception of affordances is direct—that this information is defined over our interactions with the environment—we require an explanation of what kind of physical system is capable of such behavior. While there has been some work on modeling affordances and other concepts in ecological psychology (Thill, Caligiore, Borghi, Ziemke, & Baldassarre, 2013; Zech et al., 2017; Pezzulo et al., 2011), to the best of our knowledge there are no models addressing the physical mechanism for resonance.

We agree with Raja that filling this gap in theory requires a story about the central nervous system. De-emphasizing the brain was a strategic move on the part of ecological psychologists, to redirect attention to environmental and informational constraints in the explanation of behavior. While cognitivists focused on mental representations, ecological psychologists urged the field to "Ask not what's inside your head, but what's your head inside of" (Mace, 1977). However, the field now seems ready to turn back towards the question of what's inside our heads, albeit armed with a set of theoretical desiderata for what a satisfying answer will look like.

First and foremost, an explanation of the CNS that is consistent with ecological psychology ought not to fall back on the notion of mental representations and rule-like operations. This requirement renders much of modern neuroscience as a poor foundation for ecological psychology, given the dominance of the "encodingist" view: the view that brain activity is an encoding or representation of stimulus properties, action plans, etc (Brette, 2019; Mirski & Bickhard, 2019). There have been many theoretical objections to this view, but the general thrust, as Dennett (1978) put it, is that encodingism entails an "unpaid debt of intelligence". That is, these views imply that a brain can somehow "see outside itself" to know what a pattern of its own activity represents. This debt remains unpaid because existing attempts at explanation, which may make recourse to innate, evolved knowledge structures or learning processes, run into seemingly insurmountable logical problems (Bickhard & Terveen, 1996).

Furthermore, as Anderson (2014) argues, the neuroscientific literature amassed under the encodingist assumption has ultimately undermined its own theoretical commitments. For example, while cognitivists suggested that the brain should instantiate a set of computational modules, each designed to compute a specific function, such modules have not been found. Instead, we have discovered that the brain is both highly interactive, with constant cross-talk between supposedly-distinct modules (Falandays, Batzloff, Spevack, & Spivey, 2020), and highly dynamic, with rapidly-shifting functional partnerships between brain regions constantly emerging and dissipating (M. L. Anderson, 2014; Pessoa, 2022). Seminal work from Anderson (2014) casts brain activity in terms of "transiently active local neuronal subsystems" (TALoNS), which are temporary “task-specific neural synergies that coordinate brain, body, and world” (Raja & Anderson, 2019). TALoNS have been shown to self-organize on the timescale of milliseconds in visual processing (Wu & Sabel, 2021), or on the timescale of minutes in skill acquisition (Bassett et al., 2011; Bassett, Meyer-Lindenberg, Achard, Duke, & Bullmore, 2006).

The encodingist view also finds a new challenge in recent demonstrations of the phenomenon of "representational drift," whereby supposed neural encodings change their distributed location in the brain over time (Rule, O’Leary, & Harvey, 2019). For example, O’Leary and Wyllie (2011) examined place cells in rat cortex as they repeatedly navigated a T-maze across several days. Neural recordings on day one after mastery of the T-maze showed a clear topographical mapping, but by day 10 this mapping was instantiated by an entirely different set of neurons. Similar results have been shown for odor representations in primary olfactory cortex (Schoonover, Ohashi, Axel, & Fink, 2020) and for visual representations in primary visual cortex (Deitch, Rubin, & Ziv, 2020; Marks & Goard, 2021). These findings put pressure on a representational account of neural activity, because they suggest that if neural activity is to function as a code, the brain would need to keep track of a constantly-shifting mapping from signals to meaning. Such an encoding scheme would seem rather inefficient, hence implausible from an evolutionary perspective. Furthermore, if the brain needs to track its own drifting representations, but the medium that does the tracking is subject to the same pressures, it is not clear that this would even be possible.

Ecological psychology’s solution to these challenges has been simply to abandon the search for representations. But while this may avoid the unpaid debt of intelligence, it instead incurs a debt of *resonance*. Saying what the brain is *not* won’t suffice; ecological psychology is also in need of a positive account. Towards this, Raja (2019) defines resonance as the informational coupling between two dynamical systems (Fig. 1): (1) an agent-environment system, and (2) an intra-agent system (the CNS). For dealing with the agent-environment system, ecological psychology is already equipped with established approaches. For example, Raja points to Warren's (2006) "behavioral dynamics" approach or Kelso and colleagues' "coordination dynamics" approach (Kelso, Dumas, & Tognoli, 2013), both of which describe cognition as a multi-scale dynamical system and do not appeal to computation or representation. In a similar way, Hotton and Yoshimi's (2011) “open dynamical systems” model agent-environment systems directly as dynamical systems, but also include machinery for studying the internal states that unfold in these systems. However, none of these authors commit to a specific story about the CNS.

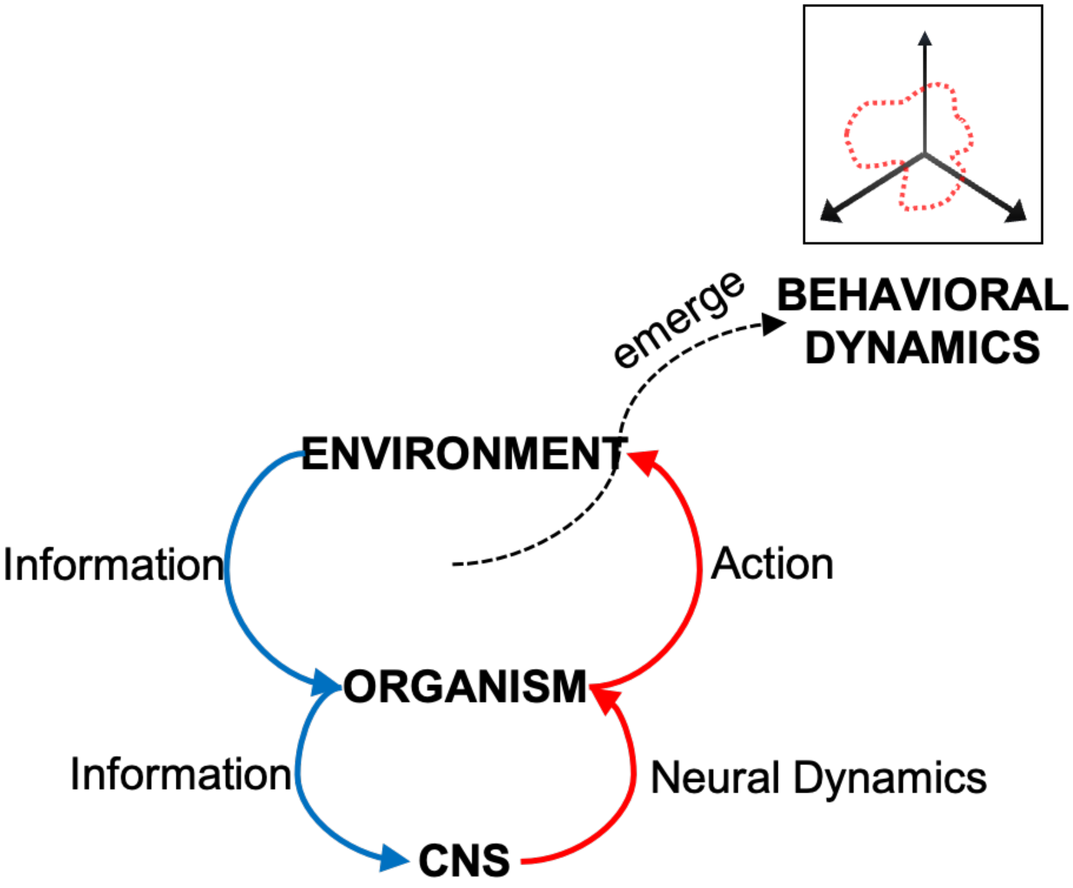

Figure 1

An illustration of the coupled CNS-Organism-Environment system. Information (blue arrows) couples the environment to the organism, and the organism to the CNS. Neural dynamics couples the CNS to the organism, and action couples the organism to the environment (red arrows). The behavioral dynamics approach (top right) focuses on emergent stabilities in the organism-environment coupling.

Davies-Barton, Raja, Baggs, and Anderson (2022) suggest autoencoders as one cognitive architecture that might be useful for ecological psychology. These are artificial neural network (ANN) architectures that learn a function to reproduce their own input. In the process, autoencoders may learn a lower-dimensional representation of the input-generating function—in other words, a model of the environment. This could allow us to preserve the idea that neural activity is an encoding *without* sneaking in any unpaid intelligence.

We agree that autoencoders have some properties that make them appealing to ecological psychology, and may indeed describe one of many possible functions implemented in the CNS. However, towards a mechanistic account of resonance, this is only a starting point, and more modeling work is needed to understand different aspects of the problem. Here, we present a complementary approach, highlight the potential utility of another artificial neural network architecture—the reservoir computer—that may be of interest to ecological psychologists due to its dynamical properties.

Reservoir computers as an ecological model of the CNS

Consider a pond of water, into which an individual throws a series of rocks at different times and locations. As the first rock is tossed in the pool, it causes a particular ripple on the pond. And as each new rock is tossed in, its own ripples interact with the radiating ripples of previous rocks. If you are a scientist, instead of having fun throwing rocks, you may stop to reflect on the fact that the state of the pond at any given moment in time—the instantaneous pattern of ripples—carries all of the information necessary to recover the locations and timings of all of the previous rocks tossed in the pond, if only you can learn to read these patterns (Yoshimi, Hotton, Tosi, & Gordon, 2022).

This is the general intuition behind both liquid state machines (Maass, Natschläger, & Markram, 2002) and echo state networks (Jaeger & Haas, 2004), introduced independently in the early 2000's, which are now grouped in the general class of algorithms called "reservoir computers." The general logic of these systems is, first, to construct a recurrent neural network without stable states, i.e. one in which activity continues to ripple through the network over time. As the network is fed a series of inputs, it will carry forward the activity from previous timesteps, therefore becoming a high-dimensional representation of the history of inputs. Researchers then need only train a simple linear readout of the reservoir state to a desired output. These networks have high computational efficiency because they only involve one layer of training a simple linear function, since the reservoir network connections typically are not adjusted, and multiple readout functions can operate in parallel on the same reservoir. Because of the integration of information across timescales, reservoir computers have been shown to be able to predict chaotic time series. And while "reservoir computer" most often refers to a class of artificial neural network models, any physical system with appropriate non-linear dynamics can play the role of a reservoir, including a literal bucket of water (Fernando & Sojakka, 2003).

Dale and Kello (2018) point out that reservoir networks are also interesting as a model of cognition, because they satisfy three important desiderata for contemporary theories of cognition. The first is dynamic memory, which refers to the fact that reservoir networks maintain a trace of past inputs in their ongoing fluctuations. This is crucial for human cognitive processes, which are clearly sensitive to contextual cues over a variety of timescales. For example, in the course of a conversation, the interpretation of a single word can be influenced by the preceding words, sentences, the entire discourse history, the identify of the speaker, and shared knowledge of events over longer timescales. Dale and

Kello point out that just *having* memory is not sufficient; memory must also be *integrated* across timescales. In reservoir networks, memory is not stored in a symbolic memory buffer, but instead embodied in the ongoing activity of the network, which allows for interaction between cues that unfold over distinct timescales, without the need to posit distinct processes for bringing together stored representations. Finally, Dale and Kello point out that reservoir networks also facilitate multimodal integration (i.e. integration of information from multiple sensory sources) in a natural way. For these reasons, they suggest that reservoir networks are particularly promising as a model of "sense-making" in human communication.

We suggest that reservoir computers have strong potential as a framework within which to model the role of CNS in action-perception more generally, and in a way that is compatible with ecological psychology. First, their oscillatory properties make them amenable to analysis within the dynamical systems framework preferred within ecological psychology, and may simply enhance biological plausibility over something like typical autoencoders[3]. Second, as Dale and Kello point out, they have several properties that make them appealing as general models of cognitive systems, including multi-modal- and multi-timescale integration. Third, and perhaps most importantly, we suggest that their activity need not be seen as representational from the perspective of the system itself, though it can be *read out* as representational to an outside observer. This is a point we will reflect on more in the next section.

But there is still one crucial way in which typical reservoir computers are unlike the CNS: they are not adaptive. In general, the weights of a reservoir network and any node properties are non-updating. However, in biological brains, change is the only constant—there is ongoing adjustment of synaptic weights, synaptogenesis or pruning, and neuron-level regulatory adjustments, among other processes. To be more useful as a model

---

[3] Note that a reservoir computer can be trained to match its own input, becoming an autoencoder, so these are not exclusive categories.

in cognitive neuroscience, reservoir computers can be amended to incorporate adaptive processes.

Self-organization in brain and behavior

An important point in ecological psychology, missing from many models of the CNS, is that brain activity and behavior are self-organizing systems (Kelso, 1995). Self-organization refers to the spontaneous emergence of structure in non-equilibrium thermodynamic systems, without the control of external agents (Prigogine & Nicolis, 1977). Consider that any human behavior, such as swinging a hammer to hit a nail, involves the coordination of many degrees of freedom (e.g. multiple limbs, joints and muscles) outside of the conscious awareness of the actor (Biryukova & Sirotkina, 2020). Somehow these many degrees of freedom constrain one another to achieve a stable outcome—hitting a nail—despite substantial variability at the lower level. In this respect, the behavioral stability can be understood as an emergent product of the interaction of many coupled degrees of freedom, without any shared representation of the goal.

It is precisely these higher-order stabilities in behavior that ecological psychology takes as its unit of analysis. Research on the self-organization of behavior highlights the functional importance of intrinsic, multiscale fluctuations (Kelty-Stephen, Palatinus, Saltzman, & Dixon, 2013; Kello et al., 2010; Pouw et al., 2021). Intrinsic fluctuations that are poised "at the edge of chaos" are thought to maximize the computational efficiency of such systems and the flexibility to switch between adaptive regimes (Bertschinger & Natschläger, 2004). Note that these properties of self-organizing systems are not specific to any level of analysis, and can apply to any system with many appropriately-coupled degrees of freedom, including the CNS. The self-organization of brain activity has become a major topic of research in its own right (Kelso, 1995; M. L. Anderson, 2014). However, in order to develop a mechanistic account of the concept of resonance, it is necessary to understand how the self-organization of CNS dynamics is linked to the self-organization of organism-environment dynamics.

An important question here is why higher-order stabilities should emerge at all in systems of many degrees of freedom, when chaos is an option. One response is to invoke the "law of maximal entropy production," which states that a "system will select the path, or assembly of paths, out of otherwise available paths, that minimize the potential or maximize the entropy at the fastest rate given the constraints" (Swenson, 1997). Under some constraints on a thermodynamic system, moving towards lower-entropy state will be the most efficient path to entropy production, and therefore order emerges spontaneously. But what exactly are the constraints that facilitate such self-organization in the CNS?

We suggest that one constraint in the CNS that may facilitate self-organization is the homeostatic tendencies of individual cells. Historically, the relevance of homeostasis to perception-action has been emphasized within the cybernetic tradition. Cyberneticists emphasized how several foundational cognitive processes can emerge from such homeostatic mechanisms. As an example of this, W. Ross Ashby (1960) offered his "homeostat", an analog computing device that adapted to maintain homeostasis in a changing environment, and in the process exhibited phenomena of learning, habituation, and reinforcement. Ecological psychology shares an emphasis on homeostasis to some extent, in that the field seeks to describe control laws for behavior. For example, work derived from Warren's (2006) behavioral dynamics approach has led to the discovery of visual control laws for locomotion in a variety of contexts (Fajen & Warren, 2007; Warren Jr & Whang, 1987; Rio, Dachner, & Warren, 2018), which often involve acting so as to cancel some change in the visual array. Thus, these laws can be thought of as mechanisms of homeostasis. Kelso's (2013) coordination dynamics approach largely uses oscillatory systems—spring equations—which is another type of homeostatic system. We suggest that a useful step towards an ecological story of the CNS is to return focus to how such organism-environment control laws may emerge from homeostatic principles in the CNS.

As O'Leary and Wyllie (2011) write, "global control is observed as an emergent feature of the nervous system, arising from the combined effects of a hierarchy of regulatory mechanisms operating on the level of cellular networks, individual cells, subcellular domains and, ultimately, individual genes and proteins" (see also: Gosak, Milojević, Duh, Skok, & Perc, 2022). This position suggests that if we build homeostatic mechanisms directly into our models of the CNS, organism-level control may emerge as a natural consequence. Neurons can regulate their own activity both by adjusting synapses as well as by modifying intrinsic properties, which may act to maintain some degree of stability in activity despite ongoing changes in the brain, such as the synaptic changes associated with Hebbian learning (Desai, 2003). For example, one important synaptic mechanism for homeostasis is heterosynaptic plasticity, by which neurons act to conserve the total weight of incoming synapses, which may help prevent runaway synaptic plasticity (Chistiakova et al., 2014; Turrigiano & Nelson, 2004). Intrinsic homeostatic mechanisms include regulating the expression of proteins that make hyperpolarizing or leak channels, which in turn may stabilize spiking frequency or resting membrane potential, for example (O'Leary & Wyllie, 2011). Although the existence of such homeostatic mechanisms are well-established, artificial neural network models do not often incorporate homeostatic principles; typically, these models focus on input-dependent synaptic adjustments (i.e. learning mechanisms such as Hebbian learning and back-propagation of errors). The few ANN models of which we are aware that have included homeostatic mechanisms (Di Paolo & Iizuka, 2008; Iizuka & Di Paolo, 2007) have used evolutionary algorithms to create viable architectures, leaving open the question of how much of their adaptability is due to homeostasis, and how much to the particular architecture that was evolved (but for a rare exception, see: Tosi, 2021).

Work by Kello and colleagues (Kello, 2013; Kello, Kerster, & Johnson, 2011; Rodny, Shea, & Kello, 2017; Szary, Kerster, & Kello, 2011) has shown that several interesting, biologically-realistic phenomena emerge when a reservoir network is endowed with homeostatic control. These researchers allowed nodes to activate or deactivate synapses in pursuit of a "critical branching ratio", meaning producing approximately one downstream spike for each of its own spikes. If this ratio is lower than 1, network activity may eventually die out, and if it is higher, activity may grow out of control—both would be bad for a human brain. The work of Kello and colleagues showed that the critical branching network produces a number of signatures of real-world self-organizing systems, including 1/f noise and neural avalanches (sudden cascades of activity with a power law distribution of magnitudes; Beggs & Plenz, 2003). Thus, in addition to having desirable properties for a cognitive architecture, reservoir networks with local homeostatic control resemble biological dynamics in important ways.

However, an important limitation of the cybernetic approach is that homeostasis was generally imposed by the researcher, who decided the homeostatic targets of the computing nodes. This approach has been criticized, from the ecological perspective, for neglecting the circular-causality in the CNS: in real biological systems, the homeostatic set point is not imposed from the outside, but instead is itself an emergent product of interaction with an environment (Turvey & Kugler, 1984). For example, in neurons, spiking activity is determined by the opening and closing of ion channels, which both influence and are influenced by the membrane potential. Due to this circular causality, the spiking activity self-stabilizes at some preferred level, which is an emergent property of the ion channel-membrane dynamics.

Falandays et al. (2021) introduced a homeostatic reservoir model that avoids this criticism to some extent, using nodes that can be described as "allostatic," meaning their homeostatic set-points are dynamic. Note that there is some debate as to whether allostatic systems are a distinct class from homeostatic systems (Corcoran & Hohwy, 2017), since homeostasis does not necessarily imply static set points (O'Leary & Wyllie, 2011), though that is often how the term has been used in practice. In the model from Falandays et al. (2021), neurons pursue homeostasis at the level of overall firing rates, while permitting of variability over time in in lower-order set points. As such, homeostasis in this model is an emergent property of the interaction of neurons with their neighbors, and in turn with the environment.

Falandays et al. (2021) suggested that neuronal homeostasis may be one potential mechanism for apparent "predictive processing" in the brain. A prominent general view in cognitive science today is that the brain learns a model of the environment by predicting upcoming sensory inputs and using prediction errors to adjust parameters of the model (Hohwy, 2018). Falandays et al. (2021) showed that some behavioral signatures associated with predictive processing can emerge from a reservoir network endowed with a neuron-level homeostatic learning rule. They presented their network with a sequence of inputs generated from a simple probabilistic grammar—four possible input "words," with a set of transitional probabilities determining the sequence. The sequence of inputs produced a sequence of perturbations across the network, which triggered homeostatic adjustments of synaptic weights and intrinsic node parameters. They found that the reservoir adapted to produce endogenous activity that compensated for the input in real time, routing inhibitory input to nodes that were receiving sensory inputs, and excitatory input to nodes that needed a boost. In other words, the network controlled its own flow of activity in a way that tracked the temporal dynamics of the input, embodying a predictive model of the input sequence for the purposes of control. As such, this model exhibited some behavioral signatures of predictive processing, such as sequence completion and spikes of activity in response to unexpected inputs, but *without* the use of explicit predictions or prediction errors.

Importantly, we believe the model introduced by Falandays et al. (2021) allows one to cast a non-representational account of how the CNS "predicts," making it potentially useful in the ecological framework. Consider that the activity in this network is not primarily an *encoding* of a current input, but instead a *complement* to the unfolding input, in the context of a dynamic neuronal milieu. Just as "one cannot step in the same river twice," as the proverb states (Graham, 2007), an input cannot perturb the network in the same way twice (at least not in practice). Because the effect of any given input on network activity may change over time, so too must the network's response to the perturbation. Along these lines, Falandays et al. (2021) found that only over relatively short timescales (dozens of input "sentences") could one discover what looked like population codes in the network—highly similar network responses to repetitions of a particular input signal—but these patterns drifted substantially over longer timescales, as the network gradually changed. Thus, although an external observer can recognize that network activity *tracks* the external input, we do not take this activity as sufficient to serve a representational function from the perspective of the network itself, since it is not associated with a stable code.

We hypothesized that local homeostatic mechanisms at the level of neurons can lead to global control at the organism level when embedded in the context of an action-perception loop. Imagine a disembodied network of neurons, with some subset that is subjected to a predictable pattern of stimulation from the environment, which produces a sequence of perturbations in spiking activity throughout the network (this describes the model in Falandays et al., 2021). If spiking activity is the variable being regulated by neurons, neuronal homeostasis and synaptic updates allow the network to eventually adapt to the regular pattern of perturbations, bringing spiking activity back towards a target profile. But consider now that, when this neural network is embodied in an organism, spiking activity may lead to movement. Movement in turn alters sensory input, leading to a different perturbation across the network. In this case, a stable signal from the environment is not a guarantee of stable input to the network, since the full input-generating process now also involves the organism's own behavior. One possible solution is to find a network state that regulates action so as to render the input regular once again[4]. For example, if the input from the environment is stable, discovering a

---

[4] What we are describing has much in common with the active-inference and free-energy minimizing approach. However, we take there to be important distinctions as well, which are beyond the scope of this article to unpack.

network state that leads to *no* motor output could render the input regular. But if the environment is itself dynamic—containing a moving stimulus, for example—then a pattern of motor output that cancels out changes in the sensory array would be a solution. In sum, we are suggesting that resonance at the organism-environment level could emerge as a stable solution to the problem of regulating activity in the CNS.

In what follows, we analyze simple models consisting of simulated mobile agents controlled by the homeostatic reservoir network introduced by Falandays et al. (2021). This work is an attempt to design a "minimally cognitive system" in the spirit of Beer et al. (1996)—a system which is as simple as possible while still producing interesting cognitive dynamics, which may help to shed light on more complex systems. We suggest that these examples illustrate a potential mechanism, at a very coarse level of description, for the Gibsonian concept of resonance in the CNS. We explore how these intrinsic fluctuations can lead agents to discover patterns of movement in a dynamic environment that serve to stabilize activity across the network—in other words, agent-environment resonance. We find that the simple homeostatic updating mechanism at the neural level spontaneously produces apparently adaptive behavior in a variety of tasks, including tracking a moving stimulus, avoiding walls, and playing the game *Pong*. Given the generalizability of this algorithm across tasks, we suggest that homeostatic reservoir networks may be an important step towards an ecological theory of the CNS.

Model Description

Network Architecture

The model consists of three layers of processing nodes: (1) an input layer, (2) a homeostatic reservoir layer, and (3) an output layer. The input represents a pattern of sensory stimulation. Input encoding is treated slightly differently in each of the cases described below. Generally speaking, nodes in this layer are tuned to spike when an input stimulus passes in front of a particular ego-centric location, analogous to light-sensitive retinal cells. Nodes in this layer do not update intrinsic parameters, and immediately reset activity at each timestep.

The input layer has non-updating feedforward projections to the reservoir layer. These links are generated randomly, with $p_{link} = .1$. Nodes within the reservoir layer also have directed connections to each other, generated randomly with $p_{link} = .1$. The connectivity matrix of the reservoir network did not update, though weights of connected nodes were allowed to change. Initial weights were randomly generated by sampling from a normal distribution with mean 0 and s.d. of 1. The reservoir nodes are discrete-time leaky integrate-and-fire nodes, which update internal parameters and incoming weights with neighbors using a homeostatic learning rule, described in the next section.

The reservoir network has feed-forward connections to an output layer, which determines a motor command. The output layer consisted of two nodes (e.g. representing a left vs. right turn motor command) with the relative strength of their activities controlling behavior. The activity of each output node is calculated as the proportion of incoming connections that propagated a spike at time t, such that output values were in the range [0,1]. Reservoir nodes were connected to output nodes with $p_{link} = .1$. Like the input layer, output nodes were non-updating and their activity was reset at each time step.

Activation Dynamics and Homeostatic Updating

The reservoir layer consists of a set of N processing nodes characterized by four intrinsic variables: (1) a current activation level x_n, initialized at 0; (2) a fixed leak rate l of .25; (3) a variable target activation level, initialized at $T_n = 1$; (4) and a variable spiking threshold $T'_n = 2T'_n$, directly coupled to target values. The value of the target T_n was given a lower bound of 1 (the value at initialization), ensuring that all nodes needed at least some continuous, positive input in order to remain near their target value. Note that targets are intrinsic parameters for each node; they do not come from an external "teaching signal."

Figure 2 shows a flowchart of the activation dynamics and homeostatic updating rules. Each iteration consists of three processing steps: (1) integrating activity, (2) spiking, and (3) homeostatic updating. In step one, nodes first leak a constant proportion l of current activation value, then sum the weighted input from external perturbations as well as from spikes within the reservoir that occurred on the previous iteration. The activation vector x of the reservoir at time step t given input vector i is:

$$x_t = (x_{t-1} \cdot l) + (i \cdot W^I) + (s(x_{t-1}) \cdot W^N) \tag{1}$$

where W^I is the input weight matrix, $s(x_{t-1})$ is the vector of length N that is equal to 1 for any node that spiked at time $t-1$ and to 0 otherwise, and W^N is the recurrent weight matrix of the reservoir.

In step two, a spike occurs when activity exceeds the spike threshold T'_n. Any node n that spikes at time t broadcasts a signal of $1 * W_{nn'}$ to connected neighbors n' at time $t+1$, while non-spiking nodes broadcast 0. The spiking node also immediately subtracts (at time t) the threshold value T'_n, producing the adjusted activation vector x'_t:

$$x'_t = x_t - s(x_t) \cdot T'_t \tag{2}$$

For example, if a node n has a current threshold $T'_n = 2$ and current activation $x^n_t = 2.5$, it will spike and drop to an activation $x'^n_t = 0.5$. Nodes can only spike once per time step, and there is no refractory period (they can spike again on the next time step).

Step three involves homeostatic updating of targets T_n (and thereby thresholds T'_n) and incoming synaptic weights. Nodes first compute the deviation from the target:

$$E^n_t = x'^n_t - T^n_t \tag{3}$$

Our homeostatic mechanism is a form of proportional control, or *P-control* in control-theory parlance, meaning that adjustments correspond to a proportion of the total error E_n. Targets were adjusted by a proportion .01 of the total error, while synaptic weights were adjusted by equally dividing the total error across all spiking neighbors.

Targets are increased if activity is above the target, or decreased if activity is below the target, unless the target is at the floor value of 1:

$$T^n_{t+1} = max(T^n_t + .01E^n_t, 0) \tag{4}$$

Incoming synaptic weights are updated in the opposite direction from targets, meaning that nodes attempt to recruit *more* input if their activity is below target, and *less* input if their activity is above target. Nodes only update weights with the subset of neighbors that spiked on the previous iteration, dividing the total error by the number of weights to be adjusted:

$$W^N_{sn,t+1} = W^N_{sn,t} - \frac{E_t}{\|W^N_{sn,t}\|} \tag{5}$$

where W^N_{sn} represents the incoming weight to node n from a neighbor s that spiked on the previous iteration, and $\|W^N_{sn,t}\|$ represents the total number of incoming weights from spiking neighbors.

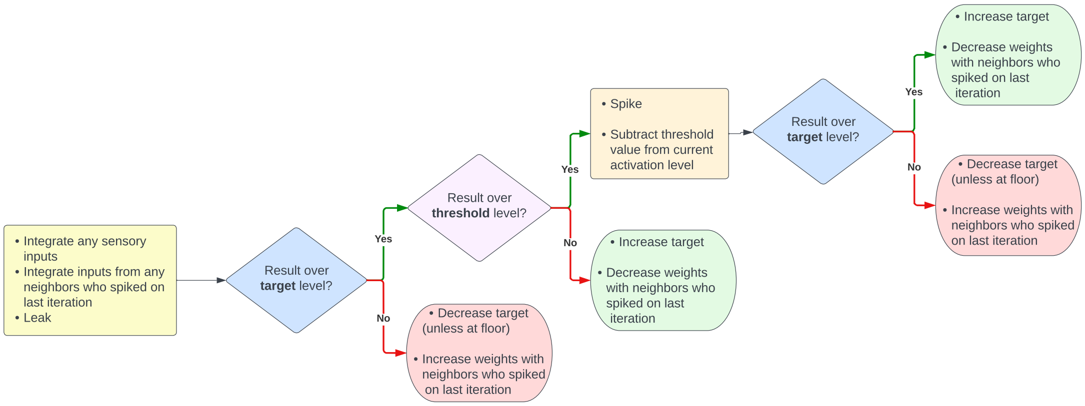

Figure 2

A flowchart displaying the homeostatic updating program. Rectangles indicate processes, diamonds indicate decision points, and rounded boxes indicate termina.

Neural Resonance and Action-Perception Loops: Three Case Studies[5]

In this subsection, we show how the homeostatic reservoir network described above may be used to control the action-perception loop of a simple agent embedded in an environment. We explore three distinct agent-environment systems: (1) an agent that can rotate right and left while position is fixed at a central point, with a stimulus that rotates around the agent at a fixed radius, (2) the classic arcade game *Pong*, where the agent corresponds to the paddle that can move up or down, with the *Pong* ball serving as stimulus, and (3) an agent similar to a Braitenberg vehicle, which can both rotate and move forward, and which senses the distance to walls in an enclosed space. In each case, inputs to the reservoir network correspond to egocentric sensory inputs based on the relative position of stimuli with respect to the agent, while the output layer controls movement. We find that with short periods of unsupervised training, the network spontaneously produces behaviors that appear adaptive in these contexts: (1) spontaneously tracking a rotating stimulus, (2) playing *Pong* with substantially above-chance performance, and (3) avoiding walls. We analyze the dynamics of the homeostatic reservoir network in the context of these agent-environment systems, showing that these adaptive behaviors are associated with drifting patterns of activity in the reservoir. These findings serve as a proof-of-concept that homeostatic mechanisms in the CNS could serve as a mechanism for agent-environment resonance, as understood in ecological psychology, while avoiding the need for a purely representational account of CNS activity.

---

[5] Data reported in this manuscript and code for running and visualizing all models is available on our Open Science Foundation repository: https://osf.io/6hqrt/. Our simulations leveraged the Agents.jl package (Datseris, Vahdati, & DuBois, 2022)

Case Study 1: Moving-Object Tracking

Agent-Environment System

The first case study, inspired by a model from Hotton and Yoshimi (2010), used the homeostatic network ($N = 200$) to control an agent that can rotate left or right while fixed at a central point. The environment contained a single stimulus that moves in a circle around the agent (see the top left panel of Fig. 3). The agent is given a set of sensor nodes that react to the presence of the stimulus, and a pair of effector nodes that allow for rotation in either direction. The stimulus moves along a circle of radius 1 at an angular speed of 1 degree per time step, thus rotating around the the agent once every 360 time steps. The simulation begins with the agent heading at 90 degrees (north) and the stimulus at 0 degrees (east), moving counter-clockwise. The stimulus was set to switch directions every 720 time steps, or two full rotations, in order to check that the agent was responsive to changing stimuli, rather than always rotating in one direction.

Sensors. The agent is imbued with 2 arrays of sensors, analogous to two eyes, positioned at +30 degrees (left sensor, red point in Fig. 3) and -30 degrees (right sensor, blue point in Fig. 3) relative to the heading angle of the agent. Each eye consists of an array of 31 input nodes (62 total for both eyes), analogous to retinal cells, that are evenly spaced in steps of 4 degrees from ±60 degrees from the center of each sensor, giving each eye a 120-degree field-of-view. Given that the left and right eyes are positioned 60 degrees apart, and that each eye contains sensors extending 60 degrees in each direction, the field-of-view for each eye overlaps in the space between them. In other words, when a stimulus is present at an angle that falls between the two eyes, both eyes are able to “see” the stimulus simultaneously.

The activity of each sensor is a Gaussian function of the angular distance of the stimulus from the respective sensor:

$$i_n = e^{-\frac{a_n^2}{10}} \tag{6}$$

where i_n is the activity of sensor n, and a is the angular distance in degrees between sensor n and the stimulus. According to this activation function, the input of sensor n was set to 1 when the stimulus was directly above the sensor, and quickly decayed to 0 when the stimulus moved further away from the sensor.

Each of the 62 input nodes was randomly connected to a node in the reservoir network with a probability of $P_{link} = .1$. The activation level of input nodes was reset at each timestep and input nodes did not utilize the homeostatic mechanism. All weights from the input to the reservoir layer were set to 0.75, and there were no connections from the reservoir to the input layer.

Effectors. In addition to having two arrays of input sensors, the agent was also given an output layer of two nodes corresponding to "effectors" for turning left or right (bottom-middle panel of Figure 3). Each node in the reservoir was randomly connected to each effector node again with a probability of $P_{link} = .1$. All connection weights from the reservoir to the output layer were set to 1.0, and there were no connections in the opposite direction. Like the input nodes, effector nodes did not use the homeostatic mechanism and their activity was reset at each timestep.

The output at each effector node was determined by the total proportion of neighbors that spiked at each time step, producing a value between 0 and 1 for each effector. For example, if an effector node had incoming connections from 20 reservoir nodes, and 10 of those reservoir nodes spiked at time t, the output of the effector was 10/20, or .5.

Movement was determined by the difference in activation value between the left- and right-turn effectors, multiplied by a gain of 10.

$$\Delta H = 10 \cdot (e_{left} - e_{right}) \tag{7}$$

where ΔH is the change in heading of the agent, and e_{left} and e_{right} represent the current output of the left and right effector nodes, respectively. Thus, if the output of $e_{left} = 1$ and $e_{right} = 0$ at time t, the agent rotated left by 10 degrees.

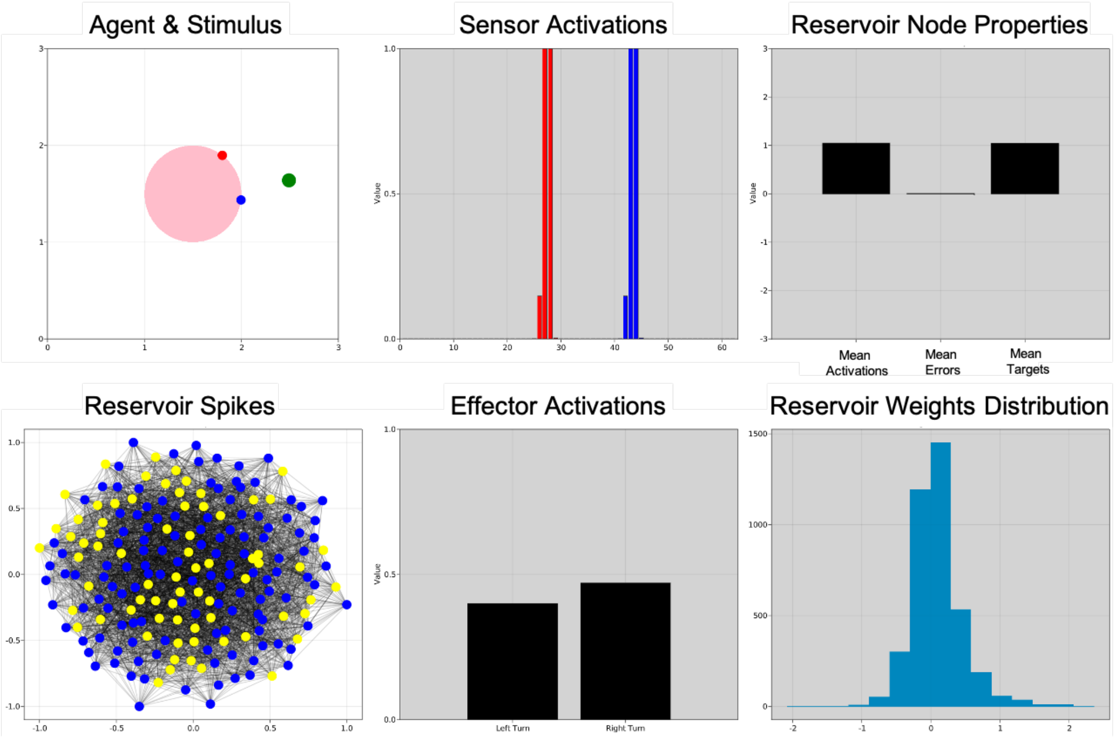

Figure 3

A still of the model as it controls the action-perception loop of a simple agent that can turn left or right. The top-left panel shows the agent (large pink circle) with two sensor arrays (centered at red and blue points) and the stimulus (green point). The top-center panel shows the activation level across the array of red and blue sensors. The top-right panel shows the current mean activation across the reservoir nodes, the mean error (discrepancy between target and activation), and mean target value. The bottom-left panel shows the reservoir, with spiking nodes shown in yellow. The bottom-middle panel shows the current activation level of the effectors for turning left (red) and right (blue). At this timestep, the stimulus is moving clockwise, and the agent is turning right (right effector > left effector) to follow it. The bottom-right panel shows the distribution of learned weights within the reservoir.

Outcomes

Spontaneous Object Tracking. When the agent's sensors first detect the presence of the stimulus, activation begins to spread through the network. This activity also spreads to the effector nodes, which initially begin moving the agent erratically left and right. After the homeostatic mechanism is applied for about 100 time steps, a sudden shift of behavior occurs: the agent locks on to the stimulus and begins rotating in the same direction, at a similar speed. When the stimulus changes directions, the agent turns to follow it with a brief delay, occasionally losing track of the stimulus. These dynamics can be seen in Figure 4, which shows the heading angle of the agent (red) and the stimulus (black) over 7200 timesteps (20 rotations of the stimulus) in a representative run.

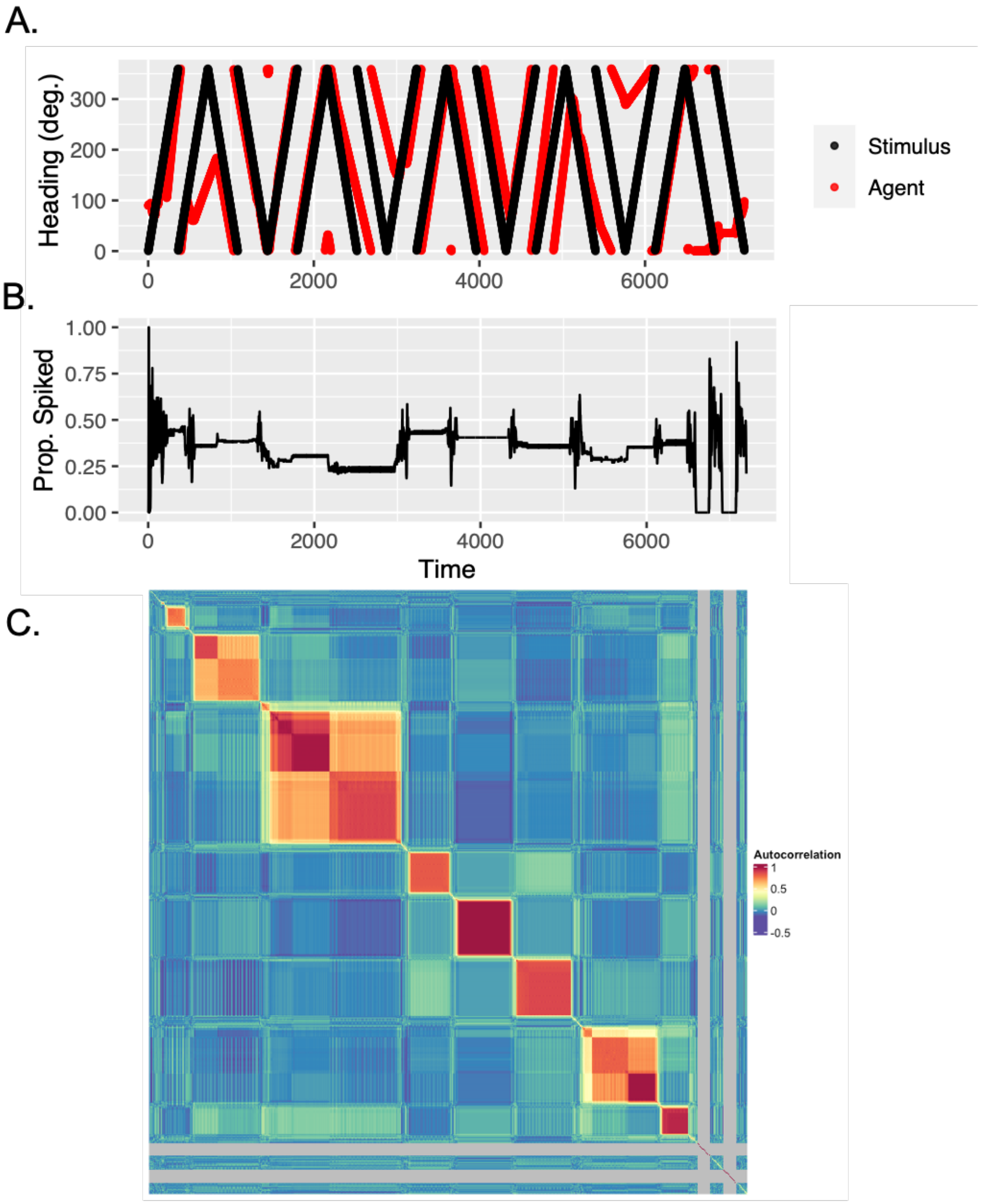

Figure 4

(A.) The heading angle of the stimulus (black lines) and the agent (red lines) over time for the first 7200 timesteps of a run. (B.) The proportion of the reservoir that was spiking at time t. (C.) The autocorrelation matrix of the spike pattern of the network for the same run. Grey bands are present for points where a correlation could not be computed because there was no variability in the spike vector (all nodes were either spiking or silent).

Why does this apparent object-tracking behavior emerge in a network that has no explicit directive to track the stimulus? This behavior can be explained by virtue of the fact that tracking the stimulus allows the network to stabilize its own activity. When the stimulus first passes over the sensors, the spikes in the network are initially chaotic. If, when this activity spreads to the effector layer, the agent turns in the *opposite* direction from the stimulus, activity will stop entering the network entirely, and the reservoir will eventually stop spiking until the stimulus comes back around (or the agent comes back around to the stimulus). Because this movement undermines the flow of input into the network, it impedes the updating of connection weights. Nodes can only update connections with neighbors that are spiking, so if the activity of the entire network dies out quickly, no updating will occur for a period of time.

On the other hand, if the activity that spreads to the effectors leads the agent to turn in the *same* direction as the stimulus, the network will continue to spike for a longer period of time, providing more opportunity for the network to learn. If the agent tracks the stimulus for a sufficient amount of time, learning can be minimized and the ongoing behavior will be sustained indefinitely. In sum, behaviors that maintain a consistent flow of input to the network are implicitly rewarded, while behaviors that undermine the input to the network are not. In this way, the network spontaneously learns to track the stimulus, “attuning” its own movements to changes in the position of the stimulus.

Transiently Active Local Neuronal Subsystems. We suggest that our model exhibits patterns reminiscent of "transiently active local neuronal subsystems" (TALoNS;

M. L. Anderson, 2014). This is most readily apparent in the autocorrelation matrix in Figure 4 C. Here, we can see that the network moves through a series of transiently-stable activity patterns (red/orange regions). Cross-referencing this figure with the proportion of the network that spiked at any given time (Fig. 4 B.) we can see that reorganizations of the network are preceded by spikes of activity. Cross-referencing again with the agent-stimulus dynamics (Fig. 4 A.), it is apparent that these spikes in activity occur either when the stimulus changes direction (e.g. around $t = 4200$), or when the agent has lost track of the stimulus and encounters it again (e.g. around $t = 3000$). At each of these time points, the agent encounters a perturbation in the flow through the network, which leads to a spike in activity that triggers homeostatic updating. This updating process results in the rapid discovery of a new local neuronal subsystem that restores stability in the network for a period of time.

Representational Drift. Next, we may also consider the degree to which this network “reuses” spike patterns over time. Given that our network appears to maintain stable tracking behavior throughout the run, except for a few brief windows where the angles of the agent and stimulus decouple, it is reasonable to expect that we may find stable patterns of activity associated with particular behavioral outcomes, such as a "turn clockwise" subnetwork and another "turn counterclockwise" subnetwork, which alternate in activity when the stimulus changes direction. However, previous work by Rodny et al. (2017) found the presence of localist representations in critical branching networks that drifted over time, similar to demonstrations of representational drift in mice (Rule et al., 2019).

As might be expected from the previous section, the autocorrelation matrix (Fig. 4 C.) suggests that any representations present must not be stable over time. Despite repeating the same behavior multiple times throughout the run, we can see that patterns associated with turning clockwise or counterclockwise at one time point are uncorrelated with patterns associated with the same behaviors at later time points. Thus, our network appears to exhibit representational drift.

Another way to visualize this representational drift is presented in Figure 5, which is similar to a figure used by (Rule et al., 2019) to show representational drift in the PPC of mice. In this figure, we plot the correlation between the spiking activity of each node, and the total output at the effector layer. Thus, strong positive correlations indicate that a particular node contributes strongly to counter-clockwise movements (higher output of effector node L), and vice versa for strong negative correlations. Each column considers these correlations in a different sliding window of 1000 timesteps, in increments of 250 timesteps. Along the diagonal, we sort the nodes in descending order according to their correlations with effector output *in that same time window*, while off-diagonal panels show nodes sorted according to a *different* time window. This plot reveals that, within any given time window we can observe what look like strong tunings for particular outcomes—particular nodes that seem to represent or encode clockwise or counter-clockwise movement. Nonetheless, when we sort nodes according to correlations in *other* time windows, we can see that these correlations fade over time. For example, nodes that were highly correlated with clockwise or counter-clockwise movement in the first time window (top-left panel) show no clear preferences for either direction during the last time window (bottom-left panel). That is, what seem like representations occur over short time windows, but these representations drift and change over time.

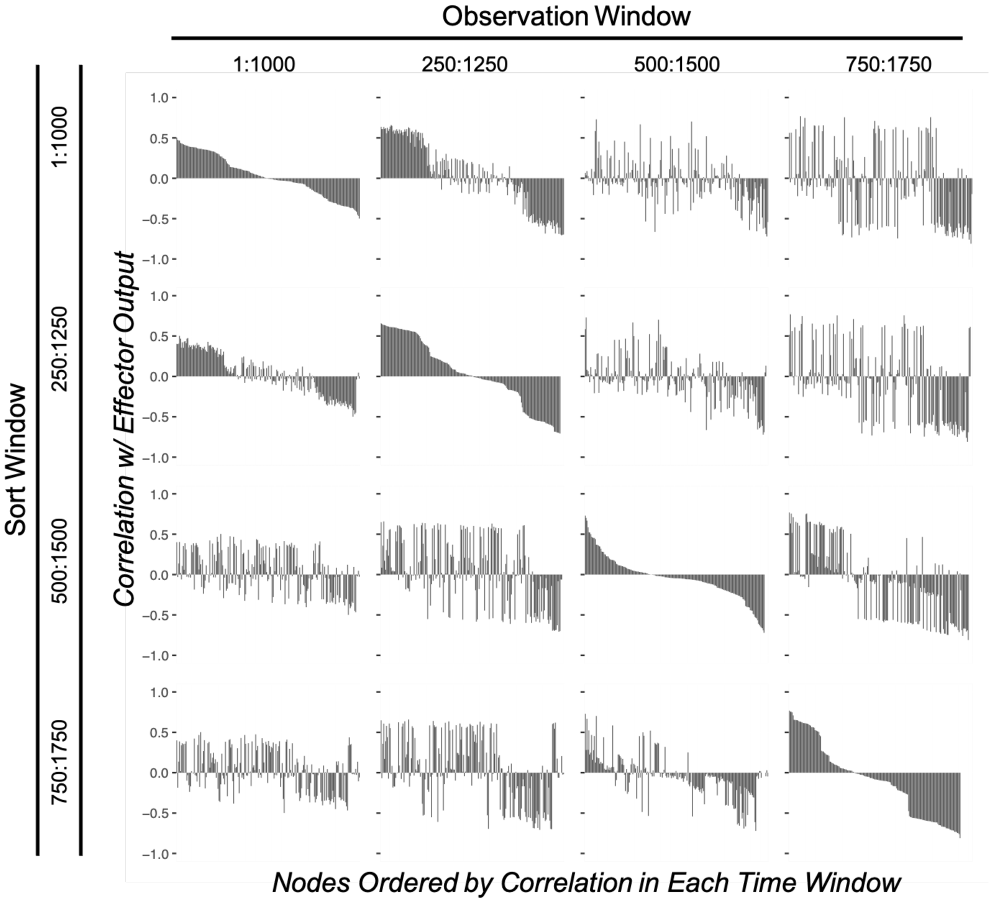

Figure 5

Representational drift in the correspondance between network activity and motor output. Each panel shows nodes on the x-axis, and their respective correlation with effector output on the y-axis. Each column corresponds to data from a sliding time window, and in each row the nodes are sorted by their correlation with effector output in a particular time window. This reveals that during any given 1000ms window, it is possible to find what appears to be a mapping between spiking activity and agent behavior, but this mapping changes substantially in as little as 250 time steps.

Case Study 2: Playing *Pong*

Background. The second case study was inspired by recent work from Kagan et al. (2022), in which a culture of cortical tissue was trained to play *Pong*. The culture was grown on a high-density microelectrode array, which received inputs based on pixel changes in the game, and generated outputs that were used to control the paddle. The culture learned to play *Pong* with slightly above-chance performance when it was trained by providing a predictable pattern of exogenous stimulation when the paddle hit the ball, or an unpredictable pattern of stimulation when the paddle missed the ball. The authors interpreted these findings through the lens of the free-energy principle (Friston, 2010a), suggesting that the cells learned to minimize prediction errors. There may be important theoretical differences between our account and those associated with the free-energy principle, which is beyond the scope of this paper to discuss in detail, but there is at least some clear overlap: homeostasis may be more achievable when patterns of stimuli are predictable, therefore a system that pursues homeostasis may act so as to render stimuli predictable.

We wondered whether our reservoir network ($N = 500$) would show similar performance in the *absence* of any exogenously-provided training signals. Note that Kagan et al. (2022) encoded sensory inputs to their cortical culture allocentrically, such that the motions of the paddle did not influence sensory input continuously (but only at discrete moments, when the paddle either hit or missed the ball, and an exogenous signal was applied). We instead coded sensory inputs egocentrically, from the perspective of the paddle. In this case, hitting the ball will naturally confer more predictable patterns of stimulation, given that a miss leads to a sudden reset of the ball's position. Unpredictable patterns of stimuli may lead to adjustments of network parameters, leading the network to search the space of parameters until a set is discovered that renders the stimulation predictable, which will consequently minimize misses. As in the first case study, tracking the movement of the ball may be implicitly rewarded by virtue of facilitating stability within the network, leading to a higher likelihood of hits.

Agent-Environment System

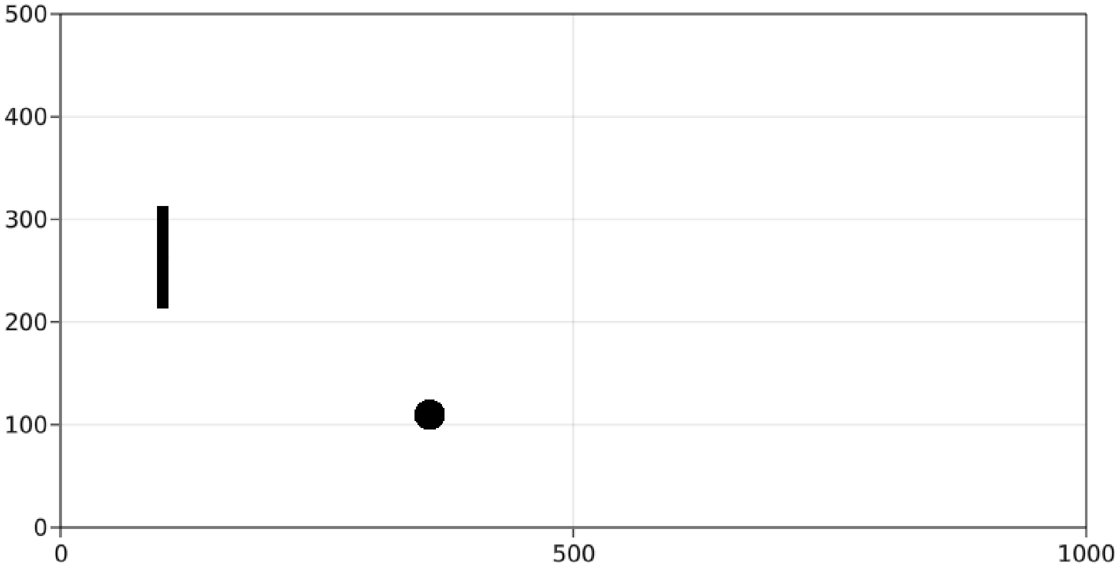

Figure 6

The game environment with paddle and ball.

The game environment consisted of a 1000 (width) X 500 (height) pixel rectangle (Fig. 6). The stimulus was the *Pong* ball, which had a radius of 15 pixels. The ball had a constant speed of 5 pixels per timestep in both the x- and y-direction. The ball was set to change y-direction upon hitting the top or bottom of the space, and to change x-direction upon hitting the right wall of the space or the paddle. The agent controlled the paddle, which was 100 pixels tall—$\frac{1}{5}$ the height of the space—making chance performance for hitting the ball 20%. The x-position of the paddle was fixed at 100, while the y-position was free to vary within the bounds of the space. If the ball passed the paddle and crossed the y-intercept ($x = 0$), the ball was immediately reset to the right side of the space with a random y-position and random y-direction.

Sensors. The agent (paddle) possessed an array of 46 sensors that radiated out from the center of the paddle over the range ±90 degrees in steps of 4 degrees. Sensors were tuned to produce an input value of 1 when angle of the stimulus (ball) relative to the center of the paddle was ≤ 2 degrees. Sensors had feed-forward connections to nodes in the reservoir layer with $P_{link} = .1$, and all input-reservoir weights were set to 2.75.

Effectors. The agent was given two effector nodes that moved the paddle up or down, with the restriction that no part of the paddle could cross the upper or lower boundary of the play area. Nodes in the reservoir network were again connected to effector nodes with $P_{link} = .1$, and motor output at each effector was taken as the proportion of spikes out of the total number of incoming connections, producing a value in the range [0, 1]. Movement was given by the relative activation of the "up" and "down" nodes, multiplied by a gain factor of 100. For example, if the "up" node was fully active and the down node was fully inactive, the paddle would move up by 100 pixels on that time step.

Outcomes

Proportion of Hits. To evaluate the success of the model in playing *Pong*, we considered the proportion of times that the paddle hit the ball out of the total number of opportunities. We ran 500 separate runs of the model, for $1x10^5$ time steps each. The mean percentage of hits over runs was 58.2% (SD = 9.95%), well above the chance performance of 20%. This data is displayed in Figure 7 panel D.

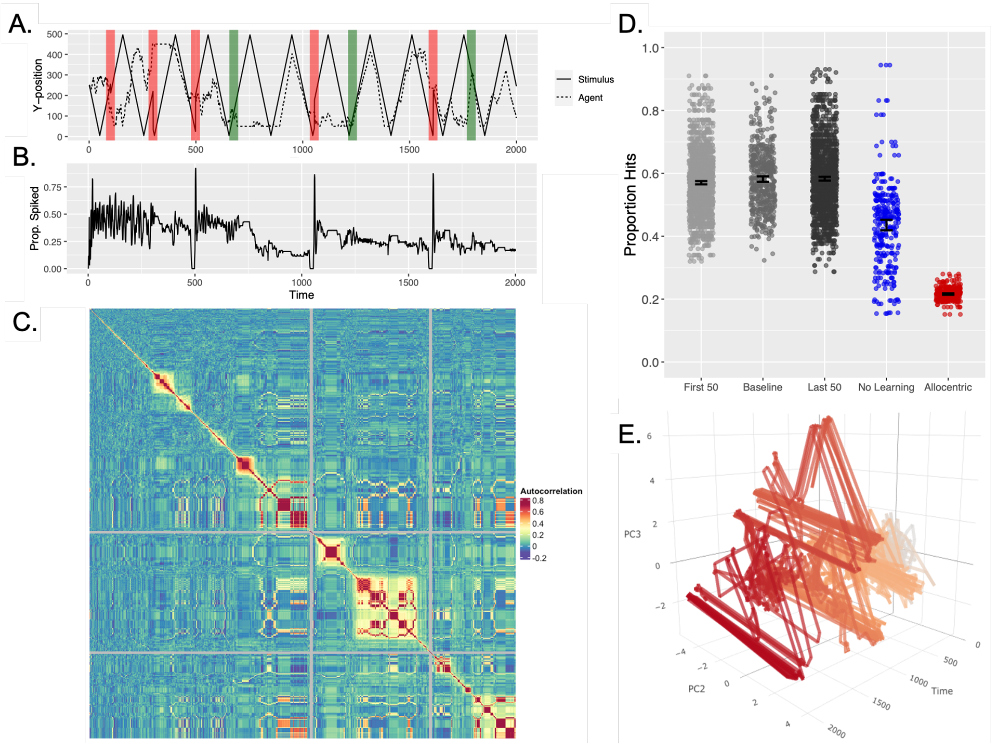

Figure 7

(A.) The y-position of the stimulus (solid lines) and the agent (dotted lines) over time for the first 2000 timesteps of a run. Red columns indicate misses, while green columns indicate hits. (B.) The proportion of the reservoir that spiked at time t on the same run. (C.) The autocorrelation matrix of the spike pattern of the network for the same run, with time moving from the top-left to the bottom-right. Warmer colors indicate higher pairwise correlation. (D.) The proportion of hits achieved in the baseline condition (medium gray), the first 50 opportunities (light gray) or last 50 opportunities in the baseline condition (dark gray), or when learning was turned off (blue), or when the sensory encoding was allocentric (red). Points correspond to individual runs, and bars display the mean and bootstrapped 95% C.I. across 500 runs. (E.) The first two principal components of the reservoir network's activity for the first 2000 timesteps of the same run. Earlier timepoints are shown in lighter colors, with later timepoints in dark red.

Learning. We next evaluated whether the model learned. One way to gauge learning is to consider whether performance improved over time. Comparing the likelihood of hits in the first 50 opportunities on each run (M = 57.86%, SD = 10.5%) to the last 50 opportunities (M = 57.86%, SD = 12%), we see that the model was at peak performance from near the beginning of a run (see Fig. 7, grayscale points). However, this should not be taken to mean that learning was irrelevant for the model's success in the task. Consider that when learning was turned off in the model, the likelihood of hits fell to 43% (SD = 13.8%; see Fig. 7, blue points). This reveals that continual adaptation of synaptic weights and internal parameters was crucial for performance.

However, it is interesting that performance with learning turned off was already well above chance performance. Why would a randomly initialized, non-updating network be inclined towards this behavior, considering there is no incentive built in to follow or hit the ball? This appears to be a natural consequence of the egocentric sensory encoding. Consider the basic law of optics that objects that are closer to an observer appear to move faster. In our model, the sensory array encodes the angle of the ball relative to the angle of the paddle, and that angle changes more rapidly as the ball approaches the paddle. Thus, the input pattern begins to change very rapidly when the ball is about to pass the paddle, which leads to an increase of activity throughout the network, which increases the likelihood of movement. Conversely, when the ball is headed towards the paddle, the angle relative to the center changes more slowly, meaning the paddle is more likely to stabilize activity and remain still. Thus, tracking the ball appears to emerge naturally from an egocentric action-perception loop in this context.

To confirm the above claims, we re-ran the model with an allocentric sensory encoding. The agent was given an array of 50 sensors that encoded the y-position of the ball in the play space, arranged in steps of 10 pixels from 5 to 495 pixels. With this encoding, the mean likelihood of hits over 100 independent runs was 21.6% (SD = 2.07%; see 7 red points), just slightly above chance level. Thus, the network only tracked the ball when sensory information was egocentric.

Case Study 3: Wall Avoidance

Background. The final case study was inspired by work from Masumori et al. (2015). Similarly to the work by Kagan et al. (2022) discussed in the previous case study, Masumori et al. (2015) grew a culture of cortical cells on a high-density microelectrode array, which was used to control a mobile robot with sensors that detected the presence of walls. They found that the collection of cells spontaneously improved in its tendency to avoid walls, without the need for any external reward. The authors proposed that this result occurs because movements that lead to the cessation of stimulation (i.e. avoidance) allow the network to stabilize activity, whereas continued stimulation leads to adaptation in the network until a stable avoidant pattern is discovered. We wondered whether our homeostatic network ($N = 200$) would produce similar results even if avoidance did *not* lead to the cessation of inputs, given our commitment to the idea that some level of continuous input is necessary for survival of neurons. Instead, we hypothesized that our network would produce movement patterns that rendered input patterns predictable, which would involve avoiding walls, given that hitting a barrier would disrupt the correspondence between motor commands and sensory inputs.

Unlike the previous two case studies, in this context the environment is entirely static. Because the full input-generating process now *only* depends upon the agent's behavior, we hypothesized that the network would be able to discover patterns of movement that rendered sensory input perfectly predictable. If this were the case, the network would eventually be able to completely stabilize its activity, and would cease changing parameters after a time. Because of this, we were also interested in how the network would perform when the input is noisy, preventing the possibility of perfect stability. Additionally, we explore the resilience of behavior to a perturbation consisting of a sudden inversion of the visual field, similar to prior work from Di Paolo and Iizuka (2008). Finally, we consider whether the network will still show some adaptive behavior when homeostatic updating was turned off, as we did in the previous case study.

Agent-Environment System

The space consisted of a simulated 15 X 15 meter box containing a circular agent of radius .5m (see Fig. 8). The agent was driven by two simulated "wheels," located ±90 degrees from the heading direction, akin to a Braitenberg vehicle, with movement driven by the relative speed of the wheels. The agent could not move any part of its circular body past a wall, and was set to suddenly rotate either +45 degrees or -45 degrees upon contacting a wall, which enhanced the degree to which hitting walls produced unpredictable patterns of stimulation.

Sensors. The agent was given two sensors, located at ±45 degrees relative to the heading direction of the agent. Each sensor casts a ray forward at the respective angle from the heading direction of the agent, and detected the nearest point of intersection with one of the four walls of the space. The strength of input at the sensor was inversely proportional to the distance to the wall, such that input was equal to 1 if the sensor was directly touching a wall, or 0 if the sensor was at the maximum distance from a wall. The maximum distance was the length of the diagonal, $\sqrt{2 \cdot 15^2}$. Because the agent could not have a sensor located perfectly in the corner of the space, as this would require having some region of the agent pass through the walls of the space, input at each sensor was always > 0. As before, sensors had feed-forward connections to nodes in the reservoir layer with $P_{link} = .1$, with all input-reservoir weights now set to 2.

Effectors. The agent had two effector nodes, which simulated motors controlling two wheels located ±90 degrees from the heading direction. As in the previous simulations, reservoir nodes were randomly connected to effectors with $P_{link} = .1$, and motor output at each effector was taken as the proportion of spikes out of the total number of incoming connections, producing a value in the range [0, 1].

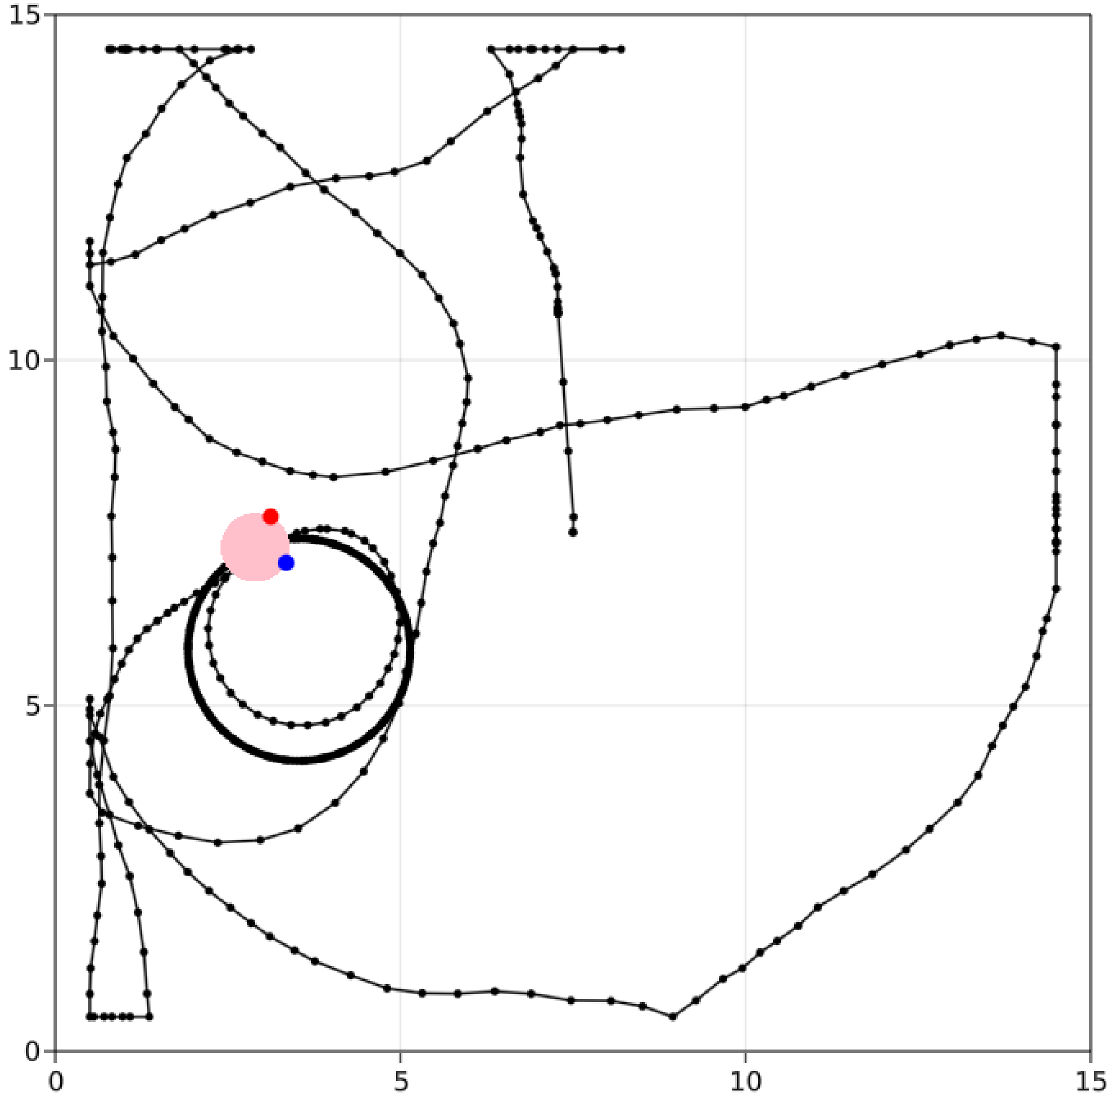

Figure 8

Outcomes

Wall Avoidance Behavior. In the absence of noise in sensory inputs, we find that the agent typically discovers a stable pattern of movement within a few hundred time steps. This stable pattern involves keeping a constant ratio of output in the left and right motors, such that the agent moves in a circle either clockwise or counter-clockwise. The ratio of outputs must be such that the circle produced has a small enough radius as to not intersect any of the boundaries of the space.

Investigation of the network dynamics once a stable movement pattern has been discovered shows that this involves either completely stable activity, or a limit cycle (repeating pattern of spikes over time) with only a small number of nodes changing values[6]. The first column of Figure 9 shows the first 1000 timesteps of a representative run in which the network discovered a stable movement pattern and network state within around 300 time steps.

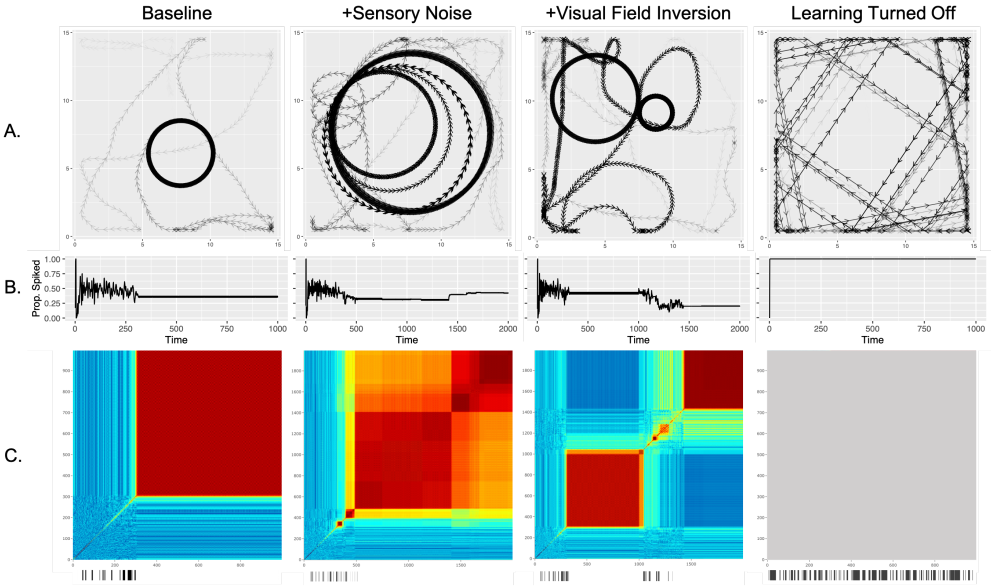

Figure 9

Row (A.) shows the trajectory of the agent in the space for the first 1000 timesteps of a representative run, with arrowheads indicating the direction of movement, and more recent points shown darker (or earlier points more transparent). Columns correspond to distinct conditions. Row (B.) shows the proportion of the reservoir that was spiking at time t*. Row (C.) shows the autocorrelation matrix of the reservoir. Notches below the x-axis indicate points at which the agent hit a wall of the space. Note that a heatmap could not be constructed for the condition with learning turned off, given that there was no variability in spiking behavior over time.*

---

[6] The behavior of this agent-environment system is similar to that of a two-vehicle Braitenberg system, studied in depth by Hotton and Yoshimi (in press). In particular, the circular behaviors are comparable to what are there studied as "revolving type relative equilibria".

Effect of Sensory Noise. Given that the static environment in this case allows the network to find a stable pattern or cycle of activity, we wondered whether the network would still show successful wall-avoidance when noise was added to sensory inputs. At each time step, we added noise by sampling from a uniform distribution in the range [-.2, .2] to each sensor independently. As such, the sensor values were now in the range [-.2, 1.2].

An example of the agent's movement dynamics from a representative run of the model under these conditions is shown in Figure 9 panel C. In the presence of noise, we find that the network can no longer completely stabilize activity. Nonetheless, the agent maintains a tendency to avoid walls, and to seek out circular movement patterns that are occasionally disrupted.

Adaptation Following Perturbation. Following work by Di Paolo and Iizuka (2008), we next examined how the network would respond to a perturbation in the form of a sudden inversion of the visual field. After 1000 time steps of the model—enough time to discover a stable movement pattern—the inputs to the left and right sensors were swapped. To amplify the perturbation caused by this change, sensor values were also multiplied by 2, such that inputs were now in the range [0,2]. An example of the agent's movement dynamics from a representative run of the model under these conditions is shown in the third column of Figure 9. Here we can see that the reservoir network takes less than 500 time steps before finding a new, completely stable pattern of activity and movement.

Learning. Finally, we considered whether the model showed evidence of learning. First, we considered the behavior of the model when homeostatic updating was turned off (right column of Fig. 9). Given that, in the case of playing *Pong*, the network shows some level of adaptive behavior even when homeostatic updating was turned off, would the same be true of the wall avoidance model? We found that in this case, when learning was turned off, activity quickly goes to a maximum, with all nodes spiking simultaneously. As a result, the model can only move straight and bounce off the walls, because the left and right effectors have equal output values.

However, when homeostatic updating is turned on, the model shows clear signs of improvement over time. In the baseline condition and the added sensory noise condition, Figure 9 shows that the agent makes a number of collisions with the wall in the first few hundred time steps (black lines beneath heatmaps), after which it discovers a stable movement pattern that results in no further contact with walls. This is also true when the visual field was inverted after 1000 time steps, except here the model again hits the wall a number of times after the inversion, before learning a new movement pattern. Thus, while learning (in the sense of improvement in performance over time) was not clearly evident in the *Pong* case study, perhaps because learning occurred too quickly to be detectable, the wall-avoidance model shows clear evidence of learning.

General Discussion

In recent years, work from Raja (2018, 2019, 2021) has called attention to the lack of a mechanistic account of the concept of "resonance" within ecological psychology, which requires a story about the CNS that does not fall back on a representationalist account of brain activity. Raja and colleagues have suggested, as a foundation for this work, Anderson’s neural reuse hypothesis, which casts brain activity in terms of transiently active local neuronal subsystems (TALoNS), which are “task-specific neural synergies that coordinate brain, body, and world” (Raja & Anderson, 2019). Adding to these arguments, we have suggested that a useful path forward is to consider the role of homeostatic properties of neurons in facilitating self-organization in the CNS. As an illustration of the utility of this view, we have considered the dynamics of simple simulated agents, endowed with minimal sensory and motor systems, mediated by a random homeostatic network. We have shown that in three distinct scenarios—(1) a rotating agent in an environment with a moving stimulus, (2) the game *Pong*, and (3) a mobile agent in a walled space—adaptive behavior spontaneously emerges. We believe that these case studies illustrate one way that the CNS could facilitate organism-environment resonance, i.e. an organism’s sensitivity to informative relations between action and perception, without relying on stabilized, computational internal representations.

What is surprising here is the fact that behavior of the model seems sensible at all right out of the box. Although work in artificial intelligence and machine-learning research has shown that similar or even more complex outcomes can be achieved with a variety of techniques, including deep neural networks with trained weights (Gibney et al., 2015), reservoir networks with a trained output mapping (Maass et al., 2002; Jaeger & Haas, 2004), networks that are evolved using a genetic algorithm (Iizuka & Di Paolo, 2007; Beer & Gallagher, 1992; Cangelosi, Parisi, & Nolfi, 1994), and of course hand-wired circuits, such as Braitenberg vehicles (Braitenberg, 1986; Hotton & Yoshimi, in press), our model has none of these features. Consider that on the first iteration of training an ANN or evolving a network, performance would typically be expected to be rather bad. So why does our model seem to exhibit reasonably context-appropriate behaviors, even exhibiting opposite patterns of behavior such as *following* a stimulus or *avoiding* walls?

The adaptive behaviors in these networks emerge spontaneously because they allow for the individual nodes to pursue homeostasis. When network activity generates movement patterns that lead to a stable flow of activity through the network, the homeostatic mechanism may reach an equilibrium, temporarily minimizing changes and therefore maintaining the ongoing behavior. The "trick" in our models is that the context-appropriate behavior just so happens to be such an equilibrium point. In scenario 1, following the stimulus keeps the sensory input stable. In scenario 2, missing the *Pong* ball leads to a sudden reset of the ball’s position, again leading to maximally unpredictable input, whereas hitting the ball preserves a continuous trajectory of inputs that changes in a predictable manner. In scenario 3, hitting a wall produces a sudden turn either clockwise or counter-clockwise. This creates a situation in which hitting walls generates maximally unpredictable flows of activity, whereas avoiding walls allows for complete stability.

Although we situate this work in the context of ecological psychology, it is worth pointing out that our interpretation of how the homeostatic mechanism leads to behavioral control is reminiscent of another theoretical paradigm, the free-energy principle (FEP; Friston, 2010b), and its related process theory, active inference (M. Ramstead, Badcock, & Friston, 2018). The FEP holds that organisms act so as to avoid "surprising" states, by optimizing a model of the environment and the consequences of actions (the "generative model"). Ecological psychology and the FEP have often been framed as competing frameworks for cognitive science (Bruineberg, Dołęga, Dewhurst, & Baltieri, 2022; though for an argument to the contrary, see Bruineberg, Kiverstein, & Rietveld, 2018), given that the former rejects representational notions, and the latter relies on them extensively. Indeed, existing FEP models of neural dynamics suggest that neural populations explicitly encode a generative model and perform Bayesian inference (M. J. Ramstead et al., 2021), and in that respect, would seem difficult to reconcile with our model. However, the FEP is framed at a higher level of abstraction than our model, and is not committed to any particular mechanistic account of brain dynamics. It is possible that an active inference model consistent with the FEP could approximate a mechanism like the one we have descrived. Purusing this question would provide a welcome opportunity for reconciliation in the “representation wars” (Constant, Clark, & Friston, 2021), but further work is needed to bear this out.

A skeptic might suggest that our examples were cherry-picked, nudged in some way to elicit the desired outcome. For example, in our wall-avoidance model, the fact that hitting a wall produces a random turn is not a necessary feature of the mechanics, and perhaps without this feature, a different pattern of behavior would emerge that would not seem "adaptive" (it should be noted that, because our model has no analog of fitness, describing the behavior as "adaptive" is purely based on our preconceived notions for what behaviors would be adaptive in a given context). While this criticism is accurate to an extent, we take these mechanics to be a reasonable analog of the real conditions faced by biological organisms. In general, some behaviors will sustain a higher-order stability in action-perception relations, while others will not. For example, while forward movement will typically produce a certain kind of optic flow, this relation will be interrupted if an organism hits a wall. If the organism was relying on a stable relation between movement and optic flow to achieve homeostasis, then hitting a wall will disrupt homeostasis, and therefore lead to a change in behavior. Thus, while our models clearly contain many simplifications, we contend that these are appropriate analogs of real constraints faced by organisms.

However, it should be noted that there are possible scenarios in which the behavior that facilitates homeostasis will *not* be the "adaptive" (per our expectations) behavior. For example, if we imagine a game of *Pong* in which the goal is to *avoid* the ball, our model would do quite poorly, because there is nothing to push it out of the observed regime of tracking the ball. In an ecological-psychology-inspired view of evolution, we would suggest that it is the role of natural selection to produce organisms for which the action-perception loops that facilitate homeostasis are precisely those that are adaptive. In other words, natural selection must generate a set of physiological constraints such that whatever flow of activity keeps *neurons* alive is also good for the entire organism. For example, if we imagine natural selection operating on a population of our *Pong*-playing agents, but in a case where *avoiding* the ball conferred fitness benefits, one possibility would be to evolve a sensory system that produces stable inputs when the ball is *not* in view, and unstable ones when the ball *is* in view. This could produce avoidance behavior without needing to search the vast space of potential networks and node types, so it may be an evolutionarily "easy" solution. Thus, while our case studies were chosen because we expected a natural correspondence between node homeostasis and adaptive behavior, we believe that these are the *typical* conditions encountered by organisms that are pre-adapted to their environments.

Furthermore, our model suggests that evolution need not act to produce highly-specific neural structures or detailed representations of the environment in order to achieve adaptive behavior. Instead, evolution needs to construct organisms that are capable of rapidly finding adaptive stabilities in the agent-environment coupling. Our examples point to a potential evolutionary "hack"—a head start on intelligence with minimal barrier to entry. Indeed, while we understand that evolution is a very important part of any theory in cognitive science, we have purposely left this out of our model at present in order to show how much intelligence can be achieved even *before* evolution has had a chance to act. Individual cells are already homeostatic, and our model shows that random collections of cells can generate behaviors that are self-preservative, context-sensitive, and rapidly adaptive to perturbations. Natural selection can then refine these simple abilities into increasingly complex behavioral repertoires simply by tuning local features of nodes (i.e. their homeostatic mechanisms, spiking mechanisms, and developmental trajectories), without needing to "know" how the entire system should act or what contexts it will encounter. This view is consistent with a developmental systems theory approach to evolution, which suggests that natural selection works upon the entire developmental trajectory, rather than simply shaping an adult phenotype (Griffiths & Tabery, 2013).

A final point worth reflecting on is our claim that the dynamics of the networks presented here need not be interpreted as "representational," as this is a crucial point for compatibility with the ecological approach[7]. In order to justify this point, we draw upon a restrictive definition of representation given by (Chemero, 2000):

> A feature R_0 of a system S will be counted as a *Representation for* S if and

---

[7] The second author, Yoshimi, interprets some points in this paper slightly differently than his co-authors, in two ways. (1) He allows that internal mediating states can be thought of as representations in a minimal sense, and has argued that that these minimal representations can have an illuminating mathematical structure, that is manifest in the "open phase portraits" of open dynamics systems (Hotton and Yoshimi, 2011, forthcoming). (2) He belives that ecological resonance of the kind described here is important and is indeed often non-representational, but also believes that an account like this can co-exist with one in which more classical forms of representation play an important role in cognitive science. In that sense he defends a form of pluralism (Yoshimi, in press)

only if:

(R1) R_0 stands between a representation producer P and a representation consumer C that have been standardized to fit one another.

(R2) R_0 has as its proper function to adapt the representation consumer C to some aspect A_0 of the environment, in particular by leading S to behave appropriately with respect to A_0, even when A_0 is not the case.

(R3) There are (in addition to R_0) transformations of R_0, R_1 ... R_n, that have as their function to adapt the representation consumer C to corresponding transformations of A_0, A_1 ... A_n

In evaluating whether our model meets these requirements for representation, first consider that, as Chemero points out, there are two distinct hypotheses regarding representation raised in the dynamical systems tradition within cognitive science. The "nature hypothesis" is an ontological claim that dynamical systems simply do not meet the requirements for using representations as outlined above, while the "knowledge hypothesis" is that some dynamical systems may still meet the requirements for using representations, but that adopting a representational stance carries no additional explanatory power. For example, Chemero notes that a paradigm example of a dynamical system, the Watt governor, actually can be given a representational description (though a non-computational or non-rule-governed version of representation), but such a view of the Watt governor is unhelpful over and above purely mathematical descriptions of its behavior.

In contrast, we suggest that our model satisfies both the nature and knowledge hypothesis, meaning that our agents neither use representations, nor can be helpfully described as such. A previous application of our model in processing a simple probabilistic grammar revealed that, in reasonably short time-windows, patterns of activation could be found that resembled population codes corresponding to specific words, or even grammatical classes that were not explicitly contained in the input data (Falandays et al., 2021). Nonetheless, these patterns drifted over time, such that *different* population codes "encoded" the same features at later time points. Similar outcomes have been observed here (see Fig. 5), with some nodes having strong correspondences with movement either clockwise or counter-clockwise, but with these correspondences drifting throughout the reservoir over time. As such, while we do see activity patterns that stand between a producer (the sensory input layer) and a consumer (the effector layer), these patterns are not stable, hence not standardized, and we can reject R1. In fact, if there are no stably recurring patterns at all, there is no entity that we can call R_0, hence nothing to ascribe functions to, nor to construct a system of representations, and therefore we can also reject R2 and R3.

In conclusion, we suggest that homeostatic networks offer a promising path towards providing a mechanistic account of the ecological concept of "resonance." This model has previously been shown to produce patterns reminiscent of predictive-processing when dealing with language (Falandays et al., 2021), and in this work has been shown capable of producing context-appropriate behavior in three distinct settings, when embedded within an action-perception loop. This suggests that our model may be applicable to an even wider variety of domains, and may help shed light on the emergence of many kinds of adaptive behavior without appeal to representation.

Acknowledgements

JBF would like to thank the following people for their helpful discussions and feedback during the development of this manuscript: Cody Moser, Paul Smaldino, Jeff Rodny, and Tim Shea from UC Merced; Charles Bakker, Noah Guzman, and Mike Anderson's EMRG lab; Daniel Friedman and the Active Inference Institute; Maxwell Ramstead and the Computational Phenomenology group; and Mac Shine.

References

Anderson, M., & Chemero, A. (2019). The world well gained. *Andy Clark and his critics*, 161–173.

Anderson, M. L. (2014). *After phrenology: Neural reuse and the interactive brain.* MIT Press.

Ashby, W. R. (1960). The homeostat. In *Design for a brain* (pp. 100–121). Springer.

Bassett, D. S., Meyer-Lindenberg, A., Achard, S., Duke, T., & Bullmore, E. (2006). Adaptive reconfiguration of fractal small-world human brain functional networks. *Proceedings of the National Academy of Sciences*, *103*(51), 19518–19523.

Bassett, D. S., Wymbs, N. F., Porter, M. A., Mucha, P. J., Carlson, J. M., & Grafton, S. T. (2011). Dynamic reconfiguration of human brain networks during learning. *Proceedings of the National Academy of Sciences*, *108*(18), 7641–7646.

Beer, R. D., & Gallagher, J. C. (1992). Evolving dynamical neural networks for adaptive behavior. *Adaptive behavior*, *1*(1), 91–122.

Beer, R. D., et al. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. *From animals to animats*, *4*, 421–429.

Beggs, J. M., & Plenz, D. (2003). Neuronal avalanches in neocortical circuits. *Journal of neuroscience*, *23*(35), 11167–11177.

Bertschinger, N., & Natschläger, T. (2004). Real-time computation at the edge of chaos in recurrent neural networks. *Neural computation*, *16*(7), 1413–1436.

Bickhard, M. H. (2016). The anticipatory brain: Two approaches. In *Fundamental issues of artificial intelligence* (pp. 261–283). Springer.

Bickhard, M. H., & Terveen, L. (1996). *Foundational issues in artificial intelligence and cognitive science: Impasse and solution* (Vol. 109). Elsevier.

Biryukova, E., & Sirotkina, I. (2020). Forward to bernstein: Movement complexity as a new frontier. *Frontiers in Neuroscience*, *14*, 553.

Braitenberg, V. (1986). *Vehicles: Experiments in synthetic psychology.* MIT press.

Brette, R. (2019). Is coding a relevant metaphor for the brain? *Behavioral and Brain Sciences*, *42*.

Bruineberg, J., Dołęga, K., Dewhurst, J., & Baltieri, M. (2022). The emperor's new markov blankets. *Behavioral and Brain Sciences*, *45*, e183.

Bruineberg, J., Kiverstein, J., & Rietveld, E. (2018). The anticipating brain is not a scientist: the free-energy principle from an ecological-enactive perspective. *Synthese*, *195*(6), 2417–2444.

Cangelosi, A., Parisi, D., & Nolfi, S. (1994). Cell division and migration in a'genotype'for neural networks. *Network: computation in neural systems*, *5*(4), 497.

Chemero, A. (2000). Anti-representationalism and the dynamical stance. *Philosophy of Science*, *67*(4), 625–647.

Chemero, A. (2003). An outline of a theory of affordances. *Ecological psychology*, *15*(2), 181–195.

Chistiakova, M., Bannon, N. M., Bazhenov, M., & Volgushev, M. (2014). Heterosynaptic plasticity: multiple mechanisms and multiple roles. *The Neuroscientist*, *20*(5), 483–498.

Constant, A., Clark, A., & Friston, K. J. (2021). Representation wars: Enacting an armistice through active inference. *Frontiers in Psychology*, *11*, 598733.

Corcoran, A. W., & Hohwy, J. (2017). Allostasis, interoception, and the free energy principle: Feeling our way forward.

Dale, R., & Kello, C. T. (2018). "how do humans make sense?" multiscale dynamics and emergent meaning. *New Ideas in Psychology*, *50*, 61–72.

Datseris, G., Vahdati, A. R., & DuBois, T. C. (2022, January). Agents.jl: a performant and feature-full agent-based modeling software of minimal code complexity. *SIMULATION*, *0*(0), 003754972110688. Retrieved from `https://doi.org/10.1177/00375497211068820` doi: 10.1177/00375497211068820

Davies-Barton, T., Raja, V., Baggs, E., & Anderson, M. L. (2022). Debt-free intelligence:

Ecological information in minds and machines.

Deitch, D., Rubin, A., & Ziv, Y. (2020). Representational drift in the mouse visual cortex. *bioRxiv*.

Dennett, D. (1978). *Philosophical essays on mind and psychology.* Bradford Books, Montgomery Vermont.

Desai, N. S. (2003). Homeostatic plasticity in the cns: synaptic and intrinsic forms. *Journal of Physiology-Paris*, *97*(4-6), 391–402.

Dewey, J. (1896). The reflex arc concept in psychology. *Psychological review*, *3*(4), 357.

de Wit, M. M., & Withagen, R. (2019). *What should a "gibsonian neuroscience" look like? introduction to the special issue* (Vol. 31) (No. 3). Taylor & Francis.

Di Paolo, E. A., & Iizuka, H. (2008). How (not) to model autonomous behaviour. *BioSystems*, *91*(2), 409–423.

Dray, W. H. (1968). On explaining how-possibly. *The Monist*, *52*(3), 390–407.

Duchon, A. P., & Warren Jr, W. H. (2002). A visual equalization strategy for locomotor control: of honeybees, robots, and humans. *Psychological Science*, *13*(3), 272–278.

Fajen, B. R., & Warren, W. H. (2007). Behavioral dynamics of intercepting a moving target. *Experimental Brain Research*, *180*(2), 303–319.

Falandays, J. B., Batzloff, B. J., Spevack, S. C., & Spivey, M. J. (2020). Interactionism in language: from neural networks to bodies to dyads. *Language, Cognition and Neuroscience*, *35*(5), 543–558.

Falandays, J. B., Nguyen, B., & Spivey, M. J. (2021). Is prediction nothing more than multi-scale pattern completion of the future? *Brain Research*, *1768*, 147578.

Fernando, C., & Sojakka, S. (2003). Pattern recognition in a bucket. In *European conference on artificial life* (pp. 588–597).

Fink, P. W., Foo, P. S., & Warren, W. H. (2009). Catching fly balls in virtual reality: A critical test of the outfielder problem. *Journal of vision*, *9*(13), 14–14.

Friston, K. (2010a). The free-energy principle: a unified brain theory? *Nature Reviews*

Neuroscience, *11*(2), 127–138.

Friston, K. (2010b). The free-energy principle: a unified brain theory? *Nature reviews neuroscience*, *11*(2), 127–138.

Gibney, E., et al. (2015). Game-playing software holds lessons for neuroscience. *Nature*, *518*(7540), 465–466.

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

Gosak, M., Milojević, M., Duh, M., Skok, K., & Perc, M. (2022). Networks behind the morphology and structural design of living systems. *Physics of Life Reviews*.

Graham, D. W. (2007). Heraclitus.

Griffiths, P. E., & Tabery, J. (2013). Developmental systems theory: What does it explain, and how does it explain it? In *Advances in child development and behavior* (Vol. 44, pp. 65–94). Elsevier.

Grossberg, S. (1982). How does a brain build a cognitive code? *Studies of mind and brain: Neural principles of learning, perception, development, cognition, and motor control*, 1–52.

Helmholtz, H. (1860). *"handbuch der physiologischen optik"* (Vol. 3). Dover.

Hohwy, J. (2018). Prediction error minimization in the brain. *The Routledge handbook of the computational mind*, 159–172.

Hotton, S., & Yoshimi, J. (2010). The dynamics of embodied cognition. *International Journal of Bifurcation and Chaos*, *20*(04), 943–972.

Hotton, S., & Yoshimi, J. (2011). Extending dynamical systems theory to model embodied cognition. *Cognitive Science*, *35*(3), 444–479.

Hotton, S., & Yoshimi, J. (in press). *The open dynamics of braitenberg vehicles.* MIT Press.

Iizuka, H., & Di Paolo, E. A. (2007). Toward spinozist robotics: Exploring the minimal dynamics of behavioral preference. *Adaptive Behavior*, *15*(4), 359–376.

Jaeger, H., & Haas, H. (2004). Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication. *science*, *304*(5667), 78–80.

Kagan, B. J., Kitchen, A. C., Tran, N. T., Habibollahi, F., Khajehnejad, M., Parker, B. J., . . . Friston, K. J. (2022). In vitro neurons learn and exhibit sentience when embodied in a simulated game-world. *Neuron*, *110*(23), 3952–3969.

Kello, C. T. (2013). Critical branching neural networks. *Psychological review*, *120*(1), 230.

Kello, C. T., Brown, G. D., Ferrer-i Cancho, R., Holden, J. G., Linkenkaer-Hansen, K., Rhodes, T., & Van Orden, G. C. (2010). Scaling laws in cognitive sciences. *Trends in cognitive sciences*, *14*(5), 223–232.

Kello, C. T., Kerster, B., & Johnson, E. (2011). Critical branching neural computation, neural avalanches, and 1/f scaling. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 33).

Kelso, J. S. (1995). *Dynamic patterns: The self-organization of brain and behavior.* MIT press.

Kelso, J. S., Dumas, G., & Tognoli, E. (2013). Outline of a general theory of behavior and brain coordination. *Neural Networks*, *37*, 120–131.

Kelty-Stephen, D. G., Palatinus, K., Saltzman, E., & Dixon, J. A. (2013). A tutorial on multifractality, cascades, and interactivity for empirical time series in ecological science. *Ecological Psychology*, *25*(1), 1–62.

Maass, W., Natschläger, T., & Markram, H. (2002). Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural computation*, *14*(11), 2531–2560.

Mace, W. M. (1977). James j. gibson's strategy for perceiving: Ask not what's inside your head, but what's your head inside of. *Perceiving, acting, and knowing: Towards an ecological psychology*.

Marks, T. D., & Goard, M. J. (2021). Stimulus-dependent representational drift in primary visual cortex. *Nature communications*, *12*(1), 1–16.

Masumori, A., Maruyama, N., Sinapayen, L., Mita, T., Frey, U., Bakkum, D., … Ikegami, T. (2015). Emergence of sense-making behavior by the stimulus avoidance principle: Experiments on a robot behavior controlled by cultured neuronal cells. In *Ecal 2015: the 13th european conference on artificial life* (pp. 373–380).

Merleau-Ponty, M. (1942). *La structure du comportement* (A. L. Fisher, Ed.). Presses Universitaires de France.

Michaels, C. F., & Palatinus, Z. (2014). A ten commandments for ecological psychology. In *The routledge handbook of embodied cognition* (pp. 19–28). Routledge.

Mirski, R., & Bickhard, M. H. (2019). Encodingism is not just a bad metaphor. *Behavioral and Brain Sciences*, *42*.

O'Leary, T., & Wyllie, D. (2011). Neuronal homeostasis: time for a change? *The Journal of Physiology*, *589*(20), 4811–4826.

Pessoa, L. (2022). The entangled brain. *Journal of Cognitive Neuroscience*, 1–12.

Pezzulo, G., Barsalou, L. W., Cangelosi, A., Fischer, M. H., McRae, K., & Spivey, M. J. (2011). The mechanics of embodiment: A dialog on embodiment and computational modeling. *Frontiers in psychology*, *2*, 5.

Pouw, W., Proksch, S., Drijvers, L., Gamba, M., Holler, J., Kello, C. T., … others (2021). Multilevel rhythms in multimodal communication.

Prigogine, I., & Nicolis, G. (1977). Self-organization. *Non-Equilibrium System*.

Raja, V. (2018). A theory of resonance: Towards an ecological cognitive architecture. *Minds and Machines*, *28*(1), 29–51.

Raja, V. (2019). From metaphor to theory: the role of resonance in perceptual learning. *Adaptive Behavior*, *27*(6), 405–421.

Raja, V. (2021). Resonance and radical embodiment. *Synthese*, *199*(1), 113–141.

Raja, V., & Anderson, M. L. (2019). Radical embodied cognitive neuroscience. *Ecological Psychology*, *31*(3), 166–181.

Ramstead, M., Badcock, P., & Friston, K. (2018). Answering schrödinger's question: A free-energy formulation. *Physics of Life Reviews*, *24*, 1–16.

Ramstead, M. J., Hesp, C., Tschantz, A., Smith, R., Constant, A., & Friston, K. (2021). Neural and phenotypic representation under the free-energy principle. *Neuroscience & Biobehavioral Reviews*, *120*, 109–122.

Rio, K. W., Dachner, G. C., & Warren, W. H. (2018). Local interactions underlying collective motion in human crowds. *Proceedings of the Royal Society B: Biological Sciences*, *285*(1878), 20180611.

Rodny, J. J., Shea, T. M., & Kello, C. T. (2017). Transient localist representations in critical branching networks. *Language, Cognition and Neuroscience*, *32*(3), 330–341.

Rule, M. E., O'Leary, T., & Harvey, C. D. (2019). Causes and consequences of representational drift. *Current opinion in neurobiology*, *58*, 141–147.

Saxberg, B. V. (1987). Projected free fall trajectories. *Biological Cybernetics*, *56*(2), 159–175.

Schoonover, C. E., Ohashi, S. N., Axel, R., & Fink, A. J. (2020). Representational drift in primary olfactory cortex. *bioRxiv*.

Srinivasan, M. V. (1992). How bees exploit optic flow: behavioural experiments and neural models. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, *337*(1281), 253–259.

Stepp, N., & Turvey, M. T. (2010). On strong anticipation. *Cognitive systems research*, *11*(2), 148–164.

Swenson, R. (1997). Autocatakinetics, evolution, and the law of maximum entropy production: A principled foundation towards the study of human ecology. *Advances in Human Ecology*, *6*, 1–48.

Szary, J., Kerster, B., & Kello, C. T. (2011). What makes a brain smart? reservoir computing as an approach for general intelligence. In *International conference on artificial general intelligence* (pp. 407–413).

Thill, S., Caligiore, D., Borghi, A. M., Ziemke, T., & Baldassarre, G. (2013). Theories and computational models of affordance and mirror systems: an integrative review. *Neuroscience & Biobehavioral Reviews*, *37*(3), 491–521.

Tosi, Z. (2021). *Self-organization in neural circuits: Proposing a new model of neural organization*. Indiana University.

Turrigiano, G., & Nelson, S. (2004). Homeostatic plasticity in the developing nervous system. *Nature Reviews Neuroscience*, *5*(2), 97–107.

Turvey, M. T., & Kugler, P. N. (1984). An ecological approach to perception and action. In *Advances in psychology* (Vol. 17, pp. 373–412). Elsevier.

Warren, W. H. (2006). The dynamics of perception and action. *Psychological review*, *113*(2), 358.

Warren Jr, W. H., & Whang, S. (1987). Visual guidance of walking through apertures: body-scaled information for affordances. *Journal of experimental psychology: human perception and performance*, *13*(3), 371.

Wu, Z., & Sabel, B. A. (2021). Spacetime in the brain: rapid brain network reorganization in visual processing and recovery. *Scientific Reports*, *11*(1), 1–12.

Yoshimi, J. (in press). Pluralist neurophenomenology: A reply to lopes. *Phenomenology and the Cognitive Sciences*.

Yoshimi, J., Hotton, S., Tosi, Z., & Gordon, C. (2022). Reservoir networks. In *Neural networks in cognitive science* (chap. 15).

Zech, P., Haller, S., Lakani, S. R., Ridge, B., Ugur, E., & Piater, J. (2017). Computational models of affordance in robotics: a taxonomy and systematic classification. *Adaptive Behavior*, *25*(5), 235–271.

Zhao, H., & Warren, W. H. (2015). On-line and model-based approaches to the visual control of action. *Vision research*, *110*, 190–202.