



Explanation Within Arm's Reach: A Predictive Processing Framework for Single Arm Use in Octopuses

Sidney Carls-Diamante¹ 

Received: 21 August 2019 / Accepted: 24 May 2021 / Published online: 1 June 2021
© The Author(s) 2021

Abstract

Octopuses are highly intelligent animals with vertebrate-like cognitive and behavioural repertoires. Despite these similarities, vertebrate-based models of cognition and behaviour cannot always be successfully applied to octopuses, due to the structural and functional characteristics that have evolved in their nervous system in response to the unique challenges posed by octopus morphology. For instance, the octopus brain does not support a *somatotopic* or point-for-point spatial map of the body—an important feature of vertebrate nervous systems. Thus, while octopuses are capable of motor tasks whose vertebrate counterparts require detailed interoceptive monitoring, these movements may not be explainable using motor control frameworks premised on internal spatial representation. One such motor task is the extension of a single arm. The ability of octopuses to select and use a single arm without the guidance of a somatotopic map has been regarded as a motor control puzzle. In an attempt at a solution, this paper develops a predictive processing account of single-arm extension in octopuses.

1 Introduction

It is somewhat of a motor control puzzle that octopuses are capable of using a single arm (Zullo and Hochner 2011)—that is, if the starting point of inquiry is a motor control model based on vertebrates, which is commonly the case in cognitive science. Cognitive science has focused largely on vertebrates for much of its history, a trend reflected in its received views and explanatory paradigms. One such historically entrenched, vertebrate-based notion is that motor control requires an internal model of the body, to be used by the motor system as a frame of reference when formulating motor commands. Because the internal model is a spatial map of the body, it enables the motor system to monitor the location of the effectors and precisely

✉ Sidney Carls-Diamante
sidney.carls-diamante@uni-konstanz.de

¹ Zukunftskolleg/Philosophy Department, Universität Konstanz, Universitätsstrasse 10, 78464 Konstanz, Germany

direct motor command volleys to specific muscles. However, such a model may not be present in octopuses, due to the inability of their central nervous system to support *somatotopic representation*, or a point-for-point map of the body (Zullo et al. 2009). Models that presuppose internal spatial representation of the body—representation-based models, for short—thus have difficulty accounting for motor control in octopuses (Gutfreund et al. 1996), in particular complex movements that supposedly require close control of the effectors. One such type of octopus behaviour is the extension of a single arm to retrieve an object.

A considerable corpus of studies—consolidated in (Levy et al. 2017) and (Zullo and Hochner 2011)—describes in detail the embodied and distributed motor control strategies octopuses use to direct their arm movements. These strategies differ radically from the representation-based motor control models familiar to cognitive science in a number of ways. One is that in octopuses, control and processing routines whose vertebrate counterparts are high-level and originating from the central nervous system (CNS) are offloaded to the peripheral nervous system (PNS). Another is the extensive use of *morphological computation*, or “the phenomenon that computation can be obtained through interactions of physical form” (Paul 2006, 619) to simplify motor control routines that would otherwise have been neurally exorbitant and beyond the capacity of the octopus’s central nervous system (Gutfreund et al. 1996). For instance, octopuses “pseudo-articulate” the arm by forming bends analogous to the joints of the human arm to accomplish tasks such as object retrieval (Sumbre et al. 2001, 2005). The location of the bend around which the arm pivots in order to ensure the correct trajectory when bringing the object closer to the body is determined by the collision point of waves of muscle activation triggered by contact with the object (Sumbre et al. 2001).

An important aspect of octopus motor control—which is the source of the apparent puzzle—is that in order to extend an arm, the higher motor centres in the brain “[generate] a single extension command that is distributed to several arms” (Zullo et al. 2009, 1633) instead of being directed towards the selected arm alone. It is believed that tactile and visual information can “filter” the extension command, activating only the selected arm while suppressing movement in the non-selected ones (Zullo et al. 2009). Nevertheless, the question remains of how these tactile and visual signals are associated with a particular arm without the guidance provided by a somatotopic map. In other words, “[if] there is no somatotopic representation of the body, how can the [octopus] determine which arm/s to move during natural behaviour?” (Zullo and Hochner 2011, 28). Due to the octopus’s lack of a somatotopic map, representation-based models are hard put to answer this question (Levy et al. 2017). In an attempt to close this explanatory gap, this paper formulates a *predictive processing* account of single-arm use in octopuses, which is premised on the notion that what the CNS transmits downstream are not explicit motor commands but predictions of proprioceptive input (Adams et al. 2013a, b).

The flow of the paper is as follows. Section 2 provides an overview of the features of the octopus nervous system. Section 3 is an exposition of representation-based motor control. Section 4 discusses predictive processing. Section 5 presents the predictive processing-based account of single arm use. Finally, Sect. 6 concludes the paper.

2 The Octopus Nervous System

The octopus nervous system is hierarchically organized and functionally decentralized, with peripheral components that are highly elaborated and autonomous. Being entirely soft-bodied with almost no hard parts, octopus morphology lacks permanent structures that can function as proprioceptive guideposts that help to simplify motor control (Levy et al. 2017). These neuroanatomical features “pose several major difficulties for accomplishing motor control that is based on the body representation scheme” (Levy et al. 2017, 3), such as the inability of the motor centres in the octopus brain to support a somatotopic or point-for-point map of the body.

With 500 million neurons—approximately the same number found in dog brains—the octopus nervous system is the largest of all invertebrate nervous systems, and is well within the vertebrate size range (Hochner 2004). These neurons are organized into the three hierarchical and autonomous components comprising the nervous system, wherein the peripheral elements function largely independently of the brain. Most of the neurons are found in the PNS, which is connected to the brain by just about 30,000 nerve fibres. These anatomical features imply that sensory information and motor commands are processed extensively in the PNS before they are transmitted upstream to the brain.

While the *brain* (CNS) is at the top of the hierarchy, it is also the smallest part of the nervous system, with 45–50 million neurons. The paired *optic lobes* are typically considered part of the PNS, but are sometimes regarded as components of the brain (Zullo and Hochner 2011). Between them, the optic lobes have 120–180 million neurons. The visual lobes are lateralized, i.e., each lobe processes the visual input from the ipsilateral eye and transmits information directly to the brain. Information from the visual lobes is then integrated in the brain.

The largest component is the peripheral *arm nervous system*,¹ with 350 million neurons distributed equally between the octopus's eight arms. A ring of fibres at the bases of the arms called the *interbrachial commissure* connects the arms to each other and to the brain. All the octopus's arms are anatomically and functionally identical (save for the third right arm in male octopuses, which is *hectocotylized* or anatomically modified for sexual function). The anatomical hierarchy within the arm nervous system, in descending order, is as follows: the axial nerve cord, five intramuscular nerve cords, sucker ganglia, and sensory receptors. The axial nerve cord is the highest-level processing and control centre in the PNS (Richter et al. 2015), and connects the brain with the lower elements in the arm nervous system, such as the muscles and suckers. It transmits high-level motor commands from the brain to the arm, consolidates sensorimotor input within the arm, and integrates central and peripheral information. The intramuscular nerve cords are responsible for the motor innervation of the arm. The rows of sucker ganglia, each one containing hundreds of neurons, innervates and

¹ A detailed description of the arm nervous system is found in P. Graziadei, "The Nervous System of the Arms," in *The Anatomy of the Nervous System of Octopus Vulgaris* (Oxford: Clarendon Press 1971), which is still considered a comprehensive reference on the subject.

receives sensorimotor input from its respective sucker. Finally, every octopus arm has about 240 million sensory receptors in the skin, suckers, and muscles, which respond to tactile, mechanical, and chemical stimuli.

The distribution of motor control labour in the octopus is likewise hierarchical (Sumbre et al. 2001, 2005, 2006). The highest motor control centres are found in the octopus brain, i.e., the basal lobe system. The brain's motor control responsibilities are "cognitive and executive functions like motor coordination, decisionmaking (sic), and learning and memory" (Levy et al. 2017, 7). For instance, the brain selects the type of movement to be executed, such as arm extension or object retrieval, and initiates the command to activate the movement. However, the "details for the execution of various motor program (sic)" (Levy et al. 2017, 8), i.e., stored patterns of information regarding the spatial parameters necessary for actualizing a particular movement, are embedded in the arm nervous system (Zullo et al. 2009, 2019). Many of these peripheral motor programs are *stereotypic*, as they specify the fixed kinematic parameters of commonly used movements (Sumbre et al. 2001, 2005).

Stimulation experiments to the basal lobes (Zullo et al. 2009)—the highest motor centres in the octopus nervous system—yielded a finding crucial to understanding motor control in octopuses: the absence of a somatotopic map to represent the octopus's body parts. This conclusion was reached because direct stimulation to the basal lobes (1) resulted in multiple arms moving identically, and (2) it was not possible to elicit the extension of a single arm. (It is ambiguous from the original text whether this refers to all of the arms, or a set of adjacent arms.) These results suggest that high-level motor commands originating from the basal lobes are transmitted globally, so that multiple arms rather than a single one receive the same extension command (Zullo et al. 2009). It was further inferred that the brain "generates only one motor command to all arms if they are activated in the same behavioral context" (Levy et al. 2017, 12).

Another significant finding was that the same type of movement can be elicited through stimulation of different areas of the basal lobes (Zullo et al. 2009). If the octopus's arms were somatotopically represented, then different areas of the basal lobes would govern different sections of the arms (see Hochner 2012). Were somatotopic representation present, stimulating a particular area of the basal lobes would produce movements varying in shape and trajectory, depending on the part of the arm they originate from. The conclusion is thus that the highest-level motor centres in the CNS are unable to proprioceptively distinguish between the octopus's individual arms, due to the absence of a somatotopic map that would have made such identification possible. Moreover, the octopus brain does not have adequate neural resources to "be able to deal with such a huge number of parameters that would be sufficient to represent its muscular system" (Levy et al. 2017, 3).

The difficulty of spatially representing the arms is further exacerbated by the absence of a rigid skeleton, which would have simplified the task by providing fixed structures that serve as proprioceptive guideposts for the motor system (Gutfreund et al. 1996; Wolpert 1997). The absence of fine-grained central mechanisms for interoceptively monitoring the body thus precludes the highest-level motor centres from formulating detailed motor commands prescribing the activation patterns for specific muscles to execute a particular movement, as is the case in vertebrates.

Instead of a spatial map of the body, the motor centres in the brain are likely to represent motor programs or functions (Zullo et al. 2009), which could be “integrated with multimodal sensory information” (Zullo and Hochner 2011, 28). Selection and activating the type of movement is thus the responsibility of the CNS. This implies that the brain registers the kind of movement in process, but may not attribute it to specific appendages, at least proprioceptively.

3 Representation-Based Motor Control

In vertebrates, representation-based motor control is the dominant paradigm. The main premise of representation-based motor control (see Wolpert 1997) is that the CNS constructs an internal spatial model of the organism from consolidated bottom-up information sourced from the exteroceptive and proprioceptive senses (Clark 2013). The model supplies the motor system with information about the position of body parts—especially the effectors—in relation to the rest of the body, as well as the structural limitations imposed by the organism’s morphology. This way, the motor system is “guided” when formulating commands that specify which muscle groups are to be activated, and how. In other words, motor control “[involves] internal representation of the target and limb positions and coordinate transformations between different internal reference frames” (Flash and Sejnowski 2001, 656).

The articulation of a model refers to the degree of accuracy and detail in which it corresponds to the anatomical segmentation and proprioceptive landmarks of the organism’s body (Grush 2001). A well-articulated model is thus of immense help to the motor system, as extensive detail makes it possible for the motor system to formulate highly specific motor command volleys. The motor command volleys are in turn directed to the appropriate muscle groups, where their activation signals are transmitted to specific muscles to produce movement.

Representation-based motor control can be described as follows. Within this paradigm, motor control is understood as the task of specifying the set of spatiotemporal coordinates and movement trajectory that bring the effectors into a particular shape (Wolpert 1997). To accomplish this task, the motor system selects a set of coordinates or degrees of freedom of movement from those that are available to the effectors concerned. The more rigid an effector is, the simpler it is to control: the movement possibilities of rigid structures are fewer than those of flexible ones, thereby narrowing down the range of degrees of freedom that are available to the motor system to choose from.

The internal model provides the motor system with information about the degrees of freedom available to the movable parts of the body. It also allows motor possibilities to be modelled offline, i.e., without the need to actually execute the desired movements (see Grush 2001, 2004). Through offline modelling, the motor system gains information about the various possible techniques in which an effector can be used to carry out a motor task, such as the different ways to hold a pencil. Each possible technique comes with its own unique set of spatiotemporal coordinates. Actualizing the movement thus entails selection and implementation of a particular technique, a process functionally equivalent to selecting one set of coordinates over

the others. Once the motor technique, i.e., a specific set of spatiotemporal coordinates, has been selected, the motor system must then transform these coordinates into motor commands that eventually produce muscle activity. In turn, motor commands work by formulating of a set of signals that activate certain muscles while suppressing others.

Representation-based motor control frameworks place a heavy emphasis on bottom-up information (Clark 2013), wherein sensorimotor signals transmitted upstream are crucial aspects of motor control. (Nevertheless, this must not be interpreted as ignoring the role of top-down information). Importantly, representation-based motor control is particularly dependent on proprioceptive information and somatotopic mapping (see Gallagher 2006). Under this view, effective motor control requires the motor system to “know” exactly to which muscle groups to direct and transmit specific activation signals. Precision and accuracy are especially important for complicated or delicate movements that demand fine-grained control in order to bring them into their target shape and position, for instance playing the violin or threading a needle. In cases such as these, the transmission of muscle activation signals and the sequence in which the muscles are to be activated must be very precise.

The importance of proprioceptive signals lies in their conveyance of information to the motor system about whether a movement is on track or not. It is not unusual for effectors to deviate from their intended trajectory, or for movements to have unexpected consequences. In such situations, the upstream proprioceptive signals (often in combination with visual feedback) that the effector transmits are used by the motor system to adjust the activation volley, thereby bringing the effector back onto the directed path, or to formulate a new volley that may be more appropriate.

Given these features of representation-based motor control, its applicability to the activation and control of single arms in octopuses is problematic. As the octopus brain does not somatotopically represent the arms (Zullo et al. 2009), it is uncertain whether a representation-based motor control schema is implemented in octopuses (Gutfreund et al. 1996; Levy et al. 2017; Zullo and Hochner 2011). Moreover, anatomical and functional studies suggest that the highest-level motor centres in the octopus brain may be incapable of formulating motor command volleys as detailed as those in vertebrates.

Instead, the motor control schema that evolved in octopuses (Hochner 2012, 2013; Sumbre et al. 2001, 2005; Zullo et al. 2009) radically differs from that of its vertebrate counterparts, on which representation-based motor control models are based. In vertebrates, the detailed spatiotemporal coordinates are prescribed in the volley of top-down motor commands that are transmitted from the brain to the effectors. In octopuses, the CNS selects the type of movement, while the arm nervous system is responsible for computing the coordinates needed to actualize it; the motor centres within the octopus brain thus do not directly activate the arm muscles.

Thus, without (1) a somatotopic map that would enable the CNS to proprioceptively distinguish between the arms (Zullo et al. 2009), and (2) direct innervation between the CNS and the arm musculature (Zullo et al. 2019), it is difficult to account for single arm extension in octopuses using a representation-based framework. However, a predictive processing framework will fare better, as will be established over the course of the next sections.

4 Predictive Processing

The basic idea behind predictive processing is that the neurocognitive system uses exteroceptive and interoceptive sensory information to construct a model of the world that is in accordance with how the organism's sensory array responds to stimuli (Clark 2016). Central to the predictive processing framework for perception and action—and therefore cognition and adaptive behaviour—is the notion of the *generative model* (not to be confused with spatial representations that internally map an organism's physical configuration).

As it does not have direct access to the environment external to the organism, the neurocognitive system depends on a hierarchical sensorimotor model—the generative model—for information about the world. This generative model is built up from sensory signals registered by the organism's sensorium. Bottom-up sensory input thus provides the neurocognitive system with the “raw facts” about what the world is like. But to be effective, the generative model also needs to “know” how these raw facts are put together, and so must account for the causal relationships that produce any given sensorimotor scenario. The generative model thus *infers* or reconstructs “how sensory signals...are generated by hidden causes” (Wiese and Metzinger 2017, 14). In reconstructing “the causal matrix responsible for...the structure” (Clark 2013, 182) of the world, the generative model in effect generates “mock sensory signals” (Wiese and Metzinger 2017, 14) that are “approximations [of the] sensory states” (Clark 2016, 94) that it registers. In inferring the causal matrix, the generative model draws on stored information or *prior beliefs*—or simply *priors*—in order to support its hypotheses about the structure of the world (Hohwy 2013). Priors are either learned through experience, or are innate, i.e., “pre-programmed” information encoded in the organism's phenotype (Ramstead et al. 2019). The generative model is thus a reconstruction of the world and its underlying causal structure based on how the organism's sensorium responds to stimuli and the information the organism already has.

One of the tasks that follows from reconstructing the causal matrix is generating *predictions*. Predictions are formed by combining priors about the world with actual sensorimotor input in order to generate hypotheses about the sensorimotor states the organism is likely to experience in the proximal or distant future, given past and present conditions. However, the generative model is susceptible to committing errors. If the actual sensorimotor state that the organism experiences conforms to the prediction, then the prediction is correct. On the other hand, if the actual sensorimotor state that is experienced differs from the prediction, then *prediction error* arises and must be corrected. Since predictions are the basis of the organism's responses to environmental conditions, it is in the organism's best interest that they be as accurate as possible. Thus, the generative model needs to be constantly fine-tuned to minimize prediction error.

Before the discussion proceeds, the issue of how a generative model may be formed in octopuses must be addressed. Evolution may have played an extensive role in the development of a generative model in octopuses, especially where

proprioception is concerned. The octopus visual system is as sophisticated as that of vertebrates. Octopuses also have an extensive capacity for short- and long-term memory, as evidenced by their behavioural plasticity and ability to learn a wide variety of tasks (Hochner et al. 2006; Mather 2008; Mather and Kuba 2013). Thus, it is well within reason to suppose that the octopus CNS is capable of storing and retrieving information or priors for use in forming predictions about exteroceptive sensory states. However, it is with regards to proprioception that the development of a generative model may differ radically from vertebrates. It is a canonical belief in cephalopod studies that the evolution of the octopus's internal monitoring and motor control mechanisms was influenced by the unique challenges posed by its morphology (Levy et al. 2017). One such evolved characteristic is the development of a repertoire of stereotypic motor programs embedded in the arm nervous system. Since these motor programs are encoded in the octopus's phenotype, the respective interoceptive sensations they generate when activated may likewise be encoded in the CNS as innate priors. If so, this would be consistent with the proposed representation of movements or motor programs, rather than body parts, in the octopus brain (Zullo et al. 2009). Consequently, upon selecting and activating a motor program, the CNS may be able to form a prediction of the interoceptive state that it will experience. Once the arm nervous system executes the movement, the signals that arise will travel upstream to the brain, where they are compared to the interoceptive prediction.

Now, back to prediction error minimization. Prediction error is minimized in two ways. The first is through *perception*, wherein the generative model is modified so that its representation of the world is made more accurate. The process of perception begins with the formulation of a prediction of such-and-such sensorimotor states, which is transmitted downstream throughout the neurocognitive system. At the same time, the sensorium registers sensory signals, which are transmitted upstream. If the actual sensorimotor signals do not match the prediction, the prediction error generated is transmitted back upstream, where it is used by the generative model to make the necessary adjustments to its reconstruction of the world. These adjustments then inform future predictions, hopefully increasing their accuracy. This process of top-down predictions being confirmed or updated by bottom-up sensory input is repeated throughout the different levels in the neurocognitive system, and until prediction error is minimized. Prediction error can also be minimized through *action*. When actual sensorimotor signals do not match the prediction, an organism can physically alter its surroundings. These changes to the environment then give rise to sensory input that conforms to the prediction.

It has also been proposed that prediction error plays an important role in motor control. The starting premise is that the higher motor control centres in the nervous system transmit top-down predictions of proprioceptive input, rather than motor commands (Adams et al. 2013a, b). A motor command is understood as “a top-down signal that activates a motor unit, i.e. a muscle,” whereas a prediction of proprioceptive input is defined as a “top-down encoding of the consequences of a movement” (Adams et al. 2013a, b, 615). As will be discussed later, a motor command contains detailed information about the spatial coordinates and trajectory that an effector must realize, as well as the muscle activation patterns necessary to execute the

movement. On the other hand, proprioceptive predictions “specify the desired consequences of a movement” (Adams et al. 2013a, b, 615), i.e., the proprioceptive and other sensory signals expected to arise when the movement is realized or completed.

Another important idea is that movement is one way by which prediction error is minimized (Adams et al. 2013a, b): proprioceptive prediction error activates motor neurons, triggering movement. Recall that action is regarded as a strategy to minimize prediction error, wherein the organism modifies the environment to make it conform to sensory predictions. Successful modification requires the control system to first determine (1) what movements are necessary to bring the surroundings into the configuration specified in the prediction and (2) the proprioceptive signals that will arise when the organism performs these movements, i.e., the *proprioceptive prediction*. Thus, the *intention*—used in the conceptually thin sense of a psychological state whose content includes a goal, which in turn specifies how the organism is to respond to a stimulus and the movements necessary to actualize that response—to execute a particular movement is formulated alongside the proprioceptive prediction (see Vance 2017). So, when moving a chair to the left is necessary for making my surroundings conform to my model of the world, my motor control system generates the intention to move the chair to the left and the prediction of the proprioceptive signals that will arise when I do so.

Importantly, minimizing prediction error requires proprioceptive predictions to be fulfilled rather than corrected (Adams et al. 2013a, b), as they specify what should be the case. When the higher motor centres formulate an intention and its accompanying proprioceptive expectations, the signals that encode expected proprioceptive input are transmitted downstream through the motor system hierarchy. However, because the effector concerned is still immobile, the proprioceptive signals it generates do not match the predictions sent from the higher levels. Prediction error is thus produced. It is this prediction error arising from the discrepancy between proprioceptive predictions and actual proprioceptive feedback that activates the motor neurons and thereby the muscles. Proprioceptive predictions are thus effectively transformed into motor commands as they travel downstream through the periphery. Because the motor units “are engaged to move [the effector] in the appropriate way, the predictions are fulfilled and prediction errors are quashed” (Vance 2017, 5).

In addition to proprioception, vision is known to play a vital role in motor control. Empirical studies have demonstrated that when proprioceptive information ceases to be available to the motor system, visual input can take over as the principal mechanism for control and monitoring (Cole and Paillard 1998; McNeill et al. 2010). Some of the most compelling findings on vision-based motor control come from studies on patients with proprioceptive *deafferentation* (e.g., McNeill et al. 2010; Aranyosi 2013), a neurological condition wherein proprioceptive signals from the body partially or completely fail to reach the central nervous system. Deafferentation usually results from interruptions in the afferent connections between the body and the brain. One of the best-documented cases of deafferentation is that of Ian Waterman (McNeill et al. 2010), who lost all proprioceptive sensation below the neck. Without interoceptive awareness of his body, Waterman had to train himself to control his movements solely through visual guidance of the effectors. Similarly, in István Aranyosi’s (2013) case, proprioceptive input and motor control of the limbs

were compromised due to temporary chemotherapy-related damage to the PNS. He reports using vision and coarse-grained proprioceptive information from proximal body parts to control his movements. Although they vary in many respects, such as the extent or aetiology of deafferentation, the cases of Waterman and Aranyosi demonstrate that motor neurons are capable of responding to visual signals.

It has thus been proposed that in the absence or interruption of proprioception, visual signals become the main source of information about the states of the effectors. Vision and other exteroceptive signals “help generate relevant prediction errors in the somatomotor hierarchy, which can help improve the accuracy of somatomotor predictions during movement” (Vance 2017, 8), proving invaluable to guiding an effector when proprioception is unavailable. To be more specific, vision-based predictions take the form of imagining the position and location of the effectors, and how they look when in motion (Vance 2017). Fulfilling such predictions thus takes place by bringing the effectors into the shape and trajectory that they are visualized to realize.

5 Prediction Error-Based Arm Activation

The motor control routine for octopus arm activation is as follows. A command to extend is generated in the higher motor centres in the brain, and is “distributed to several arms” (Zullo et al. 2009, 1633). While this command is global, it nevertheless possible for a single arm to actually extend. Zullo et al. (2009) hypothesize that this is due to a *gating* mechanism² in the octopus brain, which inhibits or suppresses responses to irrelevant input or noise (Cromwell et al. 2008). When the gating mechanism is activated, it prevents the extension command from travelling downstream in the nervous system. The gating mechanism can be released by visual or tactile signals, allowing the extension command to proceed to the selected arm or arms (Zullo et al. 2009), where the “pattern generator [confined] to the arm’s neuromuscular system” (Sumbre et al. 2001, 1846) executes and controls the movement. At the same time, these visual or tactile signals are registered by the gating mechanism as irrelevant to the non-selected arms, which are then suppressed or prevented from extending. The use of visual and tactile input to release the gating mechanism and thereby activate a specific arm is consistent with findings that octopuses tend to prefer the arm that is *directly* in the line of sight of the eye and the target object (Byrne et al. 2006a, b).

Although the gating hypothesis explains how single arm extension is possible even though the extension command is received by multiple (perhaps all) arms, there remains a gap in the account: How does the gating mechanism distinguish between

² The gating hypothesis is very recent, and is the first attempt at explaining how visual signals influence motor control in order to extend a single arm. Thus, there is still much that is unknown about the proposed gating mechanism, e.g., its location in the nervous system and how it uses visual signals to “filter” the extension command. The predictive processing account developed here is likewise an attempt to provide an answer to the latter.

the selected and non-selected arms without the guidance of a somatotopic map that would direct motor commands to the desired appendages? That is to say, how is the extension command simultaneously registered as relevant to the selected arm and irrelevant to the non-selected ones? The discussion now proceeds to answering this question, by formulating a predictive processing account of selection and activation of a single arm.

For starters, the selection of a single arm for extension is visually influenced: octopuses tend to extend the arm that lies directly in the line of sight between the eye used to view an object and the object itself (Byrne et al. 2006a). Upon seeing an object, the octopus brain then formulates what is best described as an intention. Again, an intention is defined here as a psychological state whose content includes a goal, which in turn specifies how the organism is to respond to a stimulus as well as the movements necessary to actualize that response. The content of an intention thus includes predictions of the exteroceptive and interoceptive sensory input the organism may experience upon realizing the action specified in the intention (Clark 2013). For instance, when an octopus sees a fish, its visual system transmits information to the brain, which formulates the intention to retrieve and eat the fish. A prediction then arises, specifying the sensorimotor states that arise from reaching for and grasping the fish.

This prediction consists of the (1) visual image of a particular arm reaching out and grasping the fish, and (2) the proprioceptive and tactile sensations that accompany the arm's movement as it moves to grasp the fish. The prediction is then transmitted downstream to the peripheral arm nervous system, where multiple arms receive it (Zullo et al. 2009). Because the arms are not moving at this point, the actual bottom-up sensory signals do not match the top-down prediction. Prediction error thus arises. To correct the prediction error, the peripheral neuroanatomical components responsible for actualizing the motor task (see Sumbre et al. 2001) are activated, thereby making the arm move in accordance with the prediction. Because the arm is now moving in the predicted manner, it begins to generate sensory signals that increasingly match the prediction, thereby minimizing prediction error.

Although tactile stimulation to an arm generates "robust electrical activity in the central brain" (Levy et al. 2017, 10), these signals are not plotted onto a somatotopic map. Thus, while the CNS registers information that *an* arm has been stimulated, proprioceptive identification of the specific arm is unlikely. Proprioceptive predictions generated by the brain could thus be very coarse grained: they may be able to predict both the sensation of *an* arm grasping the fish and the type of movement used to do so, but unable to identify *exactly which* arm will experience the relevant tactile and motor sensations. This account explains how centrally issued multimodal predictions are able to activate arms. But without central mechanisms of singling out a particular appendage, and multiple arms receiving the extension command, how does single arm activation take place? The visual component of the prediction appears to be the linchpin in this motor and gating task.

As with deafferented humans (Vance 2017), visual signals could be detailed enough to compensate for the coarse grain of proprioceptive predictions in octopuses. The centrally generated prediction could contain a visual image of an arm reaching for and eventually grasping the fish, which is at, say, ten o'clock relative

to the eye. Borrowing the nomenclature used by (Byrne et al. 2006b), the activated arm is designated L1, shorthand for “the first arm on the left side” of the octopus, counting from front to back. The second arm on the left hand side is referred to as L2, and so on until L4, which is the rearmost arm on the left side. The same goes for the right side, with the arms designated as R1–R4. In line with findings that visual input influences arm use (Byrne et al. 2006a), (1) the arm directly in the line of sight between the eye used and the fish is the one to be extended, and (2) the arm would follow a trajectory within this direct line of sight. Based on the documented behavioural tendencies of octopuses (Byrne et al. 2006a), as well as the actual visual input that L1 is the one directly in line with the eye and the fish, the visual component of the prediction then specifies that a particular arm, in this case L1, will handle the fish at ten o’clock. (NB: L1 is not always used in retrieval tasks, as arm selection is influenced by the parameters that determine which arm is directly in the line of sight, e.g., the eye and the angle from which the octopus looks at the fish.)

Assuming that L1 is positioned directly between the eye and the fish and that the rest of the octopus remains stationary, the visual signals that arise when the octopus looks at the fish includes information that upon activation, L1 will extend in a straight line between the eye and the fish whereas the other arms would follow a diagonal trajectory. These visual signals are then transmitted upstream from the visual lobe(s) to the brain, which formulates a prediction containing (1) the visual image of L1 reaching for the fish and (2) the proprioceptive sensation of an arm extending and eventually coming into contact with the fish. The prediction could be formed somewhere in the brain where it can be accessed by the gating mechanism that filters out the extension command in order to activate only L1.³ Since the octopus may be able to visually discriminate between its arms, as suggested by findings on the role of vision in motor control (Byrne et al. 2006a, b; Gutnick et al. 2011), the visual prediction is likely more detailed and therefore supplements the coarser-grained proprioceptive prediction.

The visual signals that the position of L1 makes it the best choice for the task then weight the visual aspect of the prediction. Furthermore, the tendency to use the arm that forms a straight line between the eye and the attended object—possibly an evolved strategy to simplify motor control—acts as a prior, adding more weight to the prediction that L1 is to be used. More attention may thus be directed towards L1 as a result of the increased weight, providing yet more weight to this aspect of the prediction (Adams et al. 2013a, b; Pareés et al. 2014). The gating mechanism, which presumably responds to visual and proprioceptive signals alike, then registers the additional weight in favour of L1 provided by the visual signals before transmitting the overall prediction down to the PNS. Moreover, in singling out L1, the visual prediction specifies the unique set of spatial coordinates relative to the fish and the eye that L1 occupies.

Because the other arms are positioned differently, they do not match the image supplied by the visual prediction that an arm will extend in a straight line to reach

³ Without detailed information about the gating mechanism, it is difficult to pinpoint where in the CNS this process occurs.

for the fish. Thus, even though activating them would fulfil the *proprioceptive* prediction that an arm is extended for fish retrieval, they would not match the *visual* prediction that the arm positioned at ten o'clock—L1—is the one to accomplish the task. A corollary of the visual prediction that L1 reaches for the fish is that the other arms are not predicted to extend. Consequently, there is no top-down visual prediction that the other arms will move, which in turn would have to be matched by bottom-up signals caused by movement in these arms in order to quash the visual prediction error. Without this prediction, there is nothing to activate the other arms (Adams et al. 2013a, b), and so the gating for the arms other than L1 is not released.

Additionally, in order to suppress movement in the non-selected arms, the gating mechanism may involve an anti-Hebbian process, wherein simultaneous activation of multiple “units in the layer...[makes] the connection between them more inhibitory, so that joint activity is discouraged” (Földiák 1990, 166). If this is true, then the extension command may be somehow “weakened” by being received by multiple arms. The combined weight supplied by the visual prediction, the prior belief that the arm that forms a straight line between the object and the eye is often used in reaching tasks, and the resulting increase in visual attention to L1 might counteract this weakening, thereby strengthening the signal enough to direct it towards and release the gating for L1. Thus, only the extension of L1 can fulfil *both* the proprioceptive and visual modalities of the prediction, and thereby correct prediction error.

An analogy to train doors with a two-step opening process may be helpful at this point. In order for a door to open, the button on it must be pressed once it turns green, signalling that the driver has enabled door opening. When a train reaches a stop, the train driver issues what can be described as a global command that allows all the doors to open, hence turning all door buttons green. In effect, this is the driver's prediction that doors will open, which in turn corresponds to the proprioceptive prediction that *an* arm will retrieve the fish. However, any given door will open only if its button is pushed. This can be compared to a passenger's prediction that a particular door will open, which then corresponds to the visual prediction that a *specific* arm is extended to grab the fish. The passenger is analogous to the gating mechanism, and her proximity to a particular door is comparable to the additional weight provided by the visual signals, priors, and attention directed to L1. Since common sense dictates that the passenger exits through the nearest door, she pushes the button on the door in front of her, in the same way that the gating mechanism allows the activation of L1.

Two predictions are therefore fulfilled when a door opens: the train driver's global prediction that doors open, and the button-pressing passenger's prediction that a particular door will open. Thus, any door that receives the general command to open *and* whose button is pushed fulfils both the driver's and the passenger's predictions, just as any arm that receives the proprioceptive prediction *and* is the subject of the visual prediction must extend in order to minimize prediction error.

However, activating an arm other than L1 is like door A opening even though it was door B's button that was pressed. Door A opening does not fulfil the passenger's prediction that her selected door B opens. In the octopus, extending another arm would not fulfil the visual prediction of L1 reaching out for the fish. Thus, activating an arm other than L1 may minimize *proprioceptive* prediction error by fulfilling the

proprioceptive prediction, but is inconsistent with the visual prediction that the arm with such-and such coordinates is to be used and so does not correct *visual* prediction error.

6 Concluding Remarks

Single arm extension via prediction error minimization can be summarized as follows. Multiple arms receive the proprioceptive prediction to extend, but only the arm that will *also* fulfil the visual prediction that specifies being directly in the line of sight between the eye and the target object is activated. In this case the appropriate arm happens to be L1, but it is not necessarily so. Thus, although the octopus brain may be unable to make a fine-grained proprioceptive prediction of which arm will be used, visual predictions could supplement this information. As such, using L1 and not the other arms generates a sensory state to the effect that “L1 is directly in the line of sight and is reaching for the fish,” as specified by the prediction. In other words, L1 is the only one that can correct prediction error by moving to fulfil the prediction.

Due to the hypothesized disparity in accuracy of proprioceptive and visual predictions, the extent of detail in signals needed to match them varies depending on the modality. Visual predictions require fine-grained confirmation, i.e., seeing exactly where the arm is, whereas proprioceptive predictions might need only the electrical signal caused by tactile stimulation (Levy et al. 2017) to indicate that physical contact between an arm and an object has been made. Nevertheless, these cross-modal signals must be consistent, complementary, and corroborative of each other.

The morphology and neurophysiology of the octopus are such that they demand accounts of motor control that differ radically from representation-based explanatory models that have traditionally dominated cognitive science. In particular, the structure and functional characteristics of the octopus nervous system, the absence of somatotopic mapping, and the use of dynamical muscle activity to bring the arms into the required shape to perform a motor task are features that make representation-based models of motor control difficult to apply to octopuses. Furthermore, there are still many unknowns about the octopus nervous system, such as the extent to which the brain uses representations of the body and if so their format and content, and whether the arm nervous system can support representation. While I have argued that a predictive processing framework effectively accounts for single-arm use in octopuses, this positive thesis does not rule out the possibility that representation-based motor control models are likewise capable of doing so—an important matter that demands an investigation of its own, and which future research would do well to look into.

Acknowledgements I am grateful to the following people: Isabella Sarto-Jackson, Nina Atanasova, Ivan Gonzalez-Cabrera, Alice Laciny, Glenn Carruthers, Monika Karmin, Emily C. Parke, Stephen Davies, and the anonymous reviewers for their comments and feedback; Luke Roelofs, Alex Gillett, Ross Pain, and Renee Jorgensen, for organizing the events wherein I presented early versions of this paper; and the audiences of the *Predictive Engines: Andy Clark and Predictive Processing* conference at Macquarie University, the *Research Colloquium on Consciousness and Cognition* at Ruhr-Universität Bochum, the

Thursday Philosophy Seminar at the Australian National University, and the *Language, Cognition, and Culture Lab* at the University of Auckland for their input on early versions of this paper. Most of the work on this paper was carried out as part of a postdoctoral fellowship at the Konrad Lorenz Institute for Evolution and Cognition Research (KLI).

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Adams, R. A., Shipp, S., & Friston, K. J. (2013a). Predictions not commands: Active inference in the motor system. *Brain Structure and Function*, *218*(3), 611–643.
- Adams, R. A., Stephan, K. E., Brown, H. R., Frith, C. D., & Friston, K. J. (2013b). The computational anatomy of psychosis. *Frontiers in Psychiatry*, *4*(47), 1–26.
- Aranyosi, I. (2013). *The peripheral mind: Philosophy of mind and the peripheral nervous system*. New York: Oxford University Press.
- Byrne, R. A., Kuba, M. J., Meisel, D. V., Griebel, U., & Mather, J. A. (2006a). Octopus arm choice is strongly influenced by eye use. *Behavioural Brain Research*, *172*, 195–201.
- Byrne, R. A., Kuba, M. J., Meisel, D. V., Griebel, U., & Mather, J. A. (2006b). Do *Octopus Vulgaris* have preferred arms? *Journal of Comparative Psychology*, *120*(3), 198–204.
- Clark, A. (2013). Whatever Next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, *36*(3), 181–253.
- Clark, A. (2016). *Surfing uncertainty: Prediction, action, and the embodied mind*. New York: Oxford University Press.
- Cole, J., & Paillard, J. (1998). Living without touch and peripheral information about body position and movement: Studeis with deafferented subjects. In Bermudez, J. (Ed.), *The body and the self* (pp. 245–266). The MIT Press.
- Cromwell, H. C., Mears, R. P., Wan, Li., & Boutros, N. N. (2008). Sensory gating: A translational effort from basic to clinical science. *Clinical EEG and Neuroscience*, *39*(2), 69–72.
- Flash, T., & Sejnowski, T. J. (2001). Computational approaches to motor control. *Current Opinion in Neurobiology*, *11*(6), 655–662.
- Földiák, P. (1990). Forming sparse representations by local anti-Hebbian learning. *Biological Cybernetics*, *64*(2), 165–170.
- Gallagher, S. (2006). *How the body shapes the mind*. Oxford: Clarendon Press.
- Graziadei, P. (1971). The nervous system of the arms. In J.Z. Young (Ed.), *The Anatomy of the nervous system of octopus vulgaris* (pp. 45–61). Oxford: Clarendon Press.
- Grush, R. (2001). The architecture of representation. In W. Bechtel, P. Mandik, J. Mundale & R. S. Stuflebeam (Eds.), *Philosophy and the neurosciences* (pp. 349–368). Blackwell Publishers Ltd.
- Grush, R. (2004). The emulation theory of representation: Motor control, imagery, and perception. *Behavioral and Brain Sciences*, *27*, 377–442.
- Gutfreund, Y., Flash, T., Yarom, Y., Fiorito, G., Segev, I., & Hochner, B. (1996). Organization of octopus arm movements: A model system for studying the control of flexible arms. *The Journal of Neuroscience*, *16*(22), 7297–7307.
- Gutnick, T., Byrne, R. A., Hochner, B., & Kuba, M. (2011). *Octopus Vulgaris* uses visual information to determine the location of its arm. *Current Biology*, *21*, 460–462.

- Hochner, B. (2004). Octopus nervous system. In G. Adelman & B. H. Smith (Eds.), *Encyclopedia of neuroscience* (3rd ed.). London: Elsevier B. V.
- Hochner, B. (2012). An embodied view of octopus neurobiology. *Current Biology*, 22, R887–R892. <https://doi.org/10.1016/j.cub.2012.09.001>
- Hochner, B. (2013). How nervous systems evolve in relation to their embodiment: What we can learn from octopuses and other molluscs. *Brain, Behavior and Evolution*, 82, 19–30.
- Hochner, B., Shomrat, T., & Fiorito, G. (2006). The octopus: A model for a comparative analysis of the evolution of learning and memory mechanisms. *Biological Bulletin*, 210(3), 308–317.
- Hohwy, J. (2013). *The Predictive mind*. Oxford: Oxford University Press.
- Levy, G., Neshet, N., Zullo, L., & Hochner, B. (2017). Motor control in soft-bodied animals: The octopus. In J. H. Byrne (Ed.), *The oxford handbook of invertebrate neurobiology* (pp. 1–27). Oxford: Oxford Handbooks Online.
- Mather, J. A. (2008). Cephalopod consciousness: Behavioural evidence. *Consciousness and Cognition*, 17, 37–48.
- Mather, J. A., & Kuba, M. J. (2013). The cephalopod specialties: Complex nervous system, learning, and cognition. *Canadian Journal of Zoology*, 91(6), 431–449.
- McNeill, D., Quaegebeur, L., & Duncan, S. (2010). IW—"the man who lost his body". In S. Gallagher & D. Schmicking (Eds.), *Handbook of phenomenology and cognitive science* (pp. 519–543). Berlin: Springer.
- Pareés, I., Brown, H., Nuruki, A., Adams, R. A., Davare, M., Bhatia, K. P., Friston, K., & Edwards, M. J. (2014). Loss of sensory attenuation in patients with functional (psycho-genic) movement disorders. *Brain*, 137, 2916–2921.
- Paul, C. (2006). Morphological computation: A basis for the analysis of morphology and control requirements. *Robotics and Autonomous Systems*, 54, 619–630.
- Ramstead, M. J. D., Constant, A., Badcock, P. B., & Friston, K. J. (2019). Variational ecology and the physics of sentient systems. *Physics of Life Reviews*, 31, 188–205.
- Richter, J. N., Hochner, B., & Kuba, M. J. (2015). Octopus arm movements under constrained conditions: Adaptation, modification and plasticity of motor primitives. *Journal of Experimental Biology*, 218, 1069–1076.
- Sumbre, G., Fiorito, G., Flash, T., & Hochner, B. (2005). Motor control of flexible octopus arms. *Nature*, 433, 595–596.
- Sumbre, G., Fiorito, G., Flash, T., & Hochner, B. (2006). Octopuses use a human-like strategy to control precise point-to-point arm movements. *Current Biology*, 16, 767–772.
- Sumbre, G., Yoram Gutfreund, G., Fiorito, T. F., & Hochner, B. (2001). Control of octopus arm extension by a peripheral motor program. *Science*, 293, 1845–1848.
- Vance, J. (2017). Action prevents error-predictive processing without active inference. In T. Metzinger & W. Wiese (Eds.), *Philosophy and predictive processing* (pp. 1–14). Frankfurt am Main: MIND Group.
- Wiese, W., & Metzinger, T. (2017). Vanilla PP for philosophers: A primer on predictive processing. In T. Metzinger & W. Wiese (Eds.), *Philosophy and predictive processing* (pp. 1–18). Frankfurt am Main: MIND Group.
- Wolpert, D. M. (1997). Computational approaches to motor control. *Trends in Cognitive Sciences*, 1(6), 209–216.
- Zullo, L., Eichenstein, H., Maiolo, F., & Hochner, B. (2019). Motor control pathways in the nervous system of octopus vulgaris arm. *Journal of Comparative Physiology A*, 205(2), 271–279.
- Zullo, L., & Hochner, B. (2011). A new perspective on the organization of an invertebrate brain. *Communicative and Integrative Biology*, 4(1), 26–29.
- Zullo, L., Sumbre, G., Agnisola, C., Flash, T., & Hochner, B. (2009). Nonsomatotopic organization of the higher motor centers in octopus. *Current Biology*, 19, 1632–1636.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.