Collaborative Project

eSMCs

Extending Sensorimotor Contingencies to Cognition

FP7-ICT-2009-6 Grant Agreement No. 270212

Deliverable D1.7 Virtual Actions and eSMCs

Mile stone:	MS4
Workpackage number:	1
Lead beneficiary number:	6 (UPV)
Estimated person months:	14
Delivery date:	48
Actual submission date:	30.03.2015
Dissemination level:	P
Contributors:	Michael Beaton, Ezequiel Di Paolo Thomas Buhrmann

Document history

Revision	Date	Modification	Authors
1	24.03.2015	Initial version	MB, EDP, TB

D1.7 Virtual Actions and eSMCs

Deliverable specification

This deliverable is linked to Task 1.5: Virtual actions: towards intention-related eSMCs

Activity type and method: A key issue regarding the extension of the SMC concept is to examine the viability of embodied theories for higher forms of cognition, in particular, cognition that is less engaged by the current sensorimotor context. In its original formulation SMCs require actions to be executed in the real world. This falls short of explaining perceptual states during paralysis or the phenomenon of synaesthesia. Here we analyse and extend the concept of SMCs to cover virtual actions. For this, we deploy the concept of virtual action developed in the study of agency in Task 1.4. We study the consistency of this concept in relation to known functional properties of neurons in frontal and prefrontal cortex. Furthermore, we investigate computationally, to what extent different non-actualised actions can be simulated in parallel in the biological substrate, or whether a sequential process is necessary. The later results would have direct implications for the relation of attention and perceptual awareness. Attention would be needed to focus on a specific action not only in the real world, but also to anticipate subsequent changes in sensory signals and hence be a necessary condition for perceptual awareness. An SMC approach to cerebellar models operating at a high resolution temporal scale (usually detached from the complete sensorimotor loop) could contribute to adding a proper modelling of automaticity in the context of hierarchical organization of eSMCs. A model integrating error signals coming from different sensory modalities as a consequence of both performed and virtual actions can provide an embodied theory of anticipatory planning that remains contextual (avoiding the frame problem).

Expected results and presentation/verification: Extension of the eSMCs theory to virtual actions; incorporation of a concept on automaticity.

Table of Contents

1. Introduction

In this report, we will discuss the relevance of virtual and counterfactual actions for sensorimotor approaches to perception and action and for experience. We define a *virtual action* as *an action that is not fully actualised, and yet has real consequences for the agent, and world.* We define a counterfactual action as an action that does not occur, but where we nevertheless have good evidence that the action would have occurred, under different circumstances. As we will see, these various species of possible action are all relevant to properly understanding and extending a sensorimotor account of cognition and experience.

Sensorimotor contingencies (SMC) theories suggest that an agent's perceptual experience depends on mastery of how their sensations change as they interact with the world. For instance, my experience of a tomato as being solid depends on my mastery, or understanding, of the fact that I could bring different parts of the tomato into view by various different motions. As various authors have pointed out (Beaton, 2013; Seth, 2014), this notion of mastery already depends on counterfactual action – what *would* happen, if I moved my head to the side, for instance? (Even if I don't in fact do so.)

In this report, we will argue that traditional approaches to action control have relied on explicit models (or internal representations) of the world. We will note that even some of the most SMC-friendly approaches (e.g. Seth, 2014) also rely on explicit internal models. We will argue that there is a possibility of critiquing these approaches, because they underplay the importance of the world itself in structuring our actions. We will then look at various architectures implemented in the eSMCs project. All of these architectures try to get away from explicitly modelling of the world. We will argue that this is a step in the right direction for modelling and understanding embodied cognition. It might seem obvious, though, that in order to master its SMCs, an agent at the least has to have an internal model of the relevant SMCs. But does this follow? Certain architectures already developed within the eSMCs project show that it is possible to master to the structure of objects in the world, without modelling those objects. In a similar manner, we will argue that, when we take account of the contribution of the world to structuring action, we can see that it is possible for an agent to master[1] SMCs without internally modelling SMCs.

We will provide a dynamical systems analysis and a few examples of what we mean by this, showing how it is possible for real, nearby (but not taken) actions to be relevant to a scientific understanding of the actual behaviour of agents. This will allow us at least to open the debate on whether any account of mastery of SMCs involving virtual actions must necessarily be internalist (i.e., implemented as models in the brain) or whether the space of plausible explanations is indeed larger and can involve accounts of the dynamic interaction between agent and world.

---

[1] In Di Paolo et al., (2014) we note that mastery in English has two senses, both of which are important for SMCs. One is meaning of mastery is the process of mastering a skill. The other meaning is being in the 'state', or situation in which one has mastered a skill. When Noë and O'Regan talk about the mastery of SMCs, they are talking about the latter: some kind of practical sensitivity to available SMCs, and their relevance for achieving one's goals. Mastery is intrinsically tied up with normativity, although this normative aspect was left implicit in the original development of SMC theory, and has been explicitly developed in Di Paolo et al., (2014) and Buhrmann et al., (2013).

Lastly, in this report, we develop some of the philosophical implications of considering a world-involving, non-internalist account, following Beaton (2013). We will argue that virtual and counterfactual actions are relevant (indeed required) in order to allow SMC theory to account for the existence of conscious experience in cases where subjects cannot act, such as locked-in syndrome. We will additionally argue that the virtual and counterfactual actions which are relevant to SMC theory are a good objective match for first-person phenomenology. That is, we will argue that the structure of actual (and potential) world-involving dynamics are a better match for first-person experience than simply brain-involving dynamics. This is contrary to what is normally supposed, even (we will argue) within the traditionally enactive literature on neurophenomenology (Varela, 1996; Gallagher and Zahavi, 2008).

2. Approaches to modelling virtual action

In terms of modelling, the notion of virtual action inherits at least some of the explanatory burden of what in classical AI goes under the name of *planning*. Typically, whether in real robotic systems or in other contexts, planning in GOFAI involved always maintaining and using an accurate enough representation of the world. This form of strong representationalism was characterized by explicit representations of objects, contexts and of the actions that could be performed, usually at a very high level of description (e.g., "pick-up", "move", "drop", and so on).

For our purposes it is more relevant to examine in a bit more detail what we might call action-based approaches; that is, approaches which have already moved on, to a greater or lesser extent, from the constraints of traditional AI. Here the use of representations is not necessarily dropped altogether, but their use involve progressively more ego-centric, bodily information; in some cases, this ego-centric information is cast explicitly in the form of sensorimotor contingencies.

2. 1. Existing non-project approaches

Internal models

Computational theories of motor control argue for the existence and even the necessity of sophisticated inverse or forward or models of the dynamics of the human body implemented in neural circuits (Wolpert and Miall, 1996; Kawato, 1999; Ostry and Feldman, 2003). Such models, it is claimed, are crucial for multi-joint motor control. According to these theories, the problem of movement generation is essentially one of coordinate transformations: a desired movement is first defined in spatial coordinates and then transformed into the required forces or torques to be applied at the joints by the muscles. To compute these variables, the system uses internal representations of the geometry as well as the dynamical equations of motion of the body. The former can be used to calculate inverse kinematics, e.g. the joint angles required to position the hand at a particular point in external Cartesian space. The latter is used in an inverse dynamics process to calculate the joint torques or muscle forces necessary to drive to those desired joint angles.

We will not rehearse in this report the criticisms of these computational theories or discuss alternatives (this is extensively discussed in Deliverables 1.1 and 1.2)[2]. We will focus instead on how these and similar ideas have informed notions akin to that of virtual actions. It is clear that from a computational perspective the availability of internal models of sensorimotor activity could be beneficial for trajectory planning, without involving necessarily the actual execution of such trajectories (Kawato, 1999). This accounts for the popularity of internal models of various kinds for explaining complex cognition. For instance, Grush, in his emulation theory (2004), proposes that internal models of the body and the environment, whose primary purpose is to improve motor control and perception, can be run offline to produce motor and visual imagery, help interpret theory-of-mind phenomena, and could ground the semantics of language in their behavior-oriented representational structure (see also Clark and Grush, 1999). Wolpert et al. (2003) have also argued that internal models can be used in action understanding, imitation, theory of mind, and social interaction. There is a growing trend arguing that organisms need to be able to predict the sensory consequences of their own actions to be capable of fast and adaptive behavior, and that internal modeling provides a unifying framework for understanding cognition in general.

Unlike more classical GOFAI representational approaches, internal models seem more amenable to an SMCs approach since they are based on modelling internal SMCs (e.g., mappings and transformations between neural activity, torques and proprioception). And even if for this reason they could be considered largely internalist, they can easily be adapted to representational theories that involve features of the coupling with the world. Wheeler (2005) indeed argues that one could take into account the criticisms of embodied-embedded cognition and still rescue a notion of action-oriented representations. Following Clark (1997) he agrees that through such representations, objects in the world would themselves be "*encoded in terms of* possibilities for action" (p. 197). In other words, such representations would be ego-centric, contextual and oriented towards action possibilities, they would represent objects as a bundle of actions that are not necessarily taken, i.e., as a virtual action according to our definition.

Another fashionable version of internal models goes under the label of predictive coding or predictive processing (e.g., Clark 2013, Friston, 2008). In the tradition of Powers' (1973) perceptual control theory, predictive processing models are organized hierarchically and provide top-down processing of the expectation of action outcomes. Accordingly, to perceive the world is to successfully predict our own sensory states. This is done by the brain using stored knowledge about the world and the probabilities of one state or event following another to generate a prediction of what the current state is likely to be, given the previous one and this body of knowledge. Mismatches between the prediction and the received signal generate error signals that can drive learning. Unlike perceptual control theory (Powers, 1973), the top down predictions are implemented in hierarchical generative models expressed in probability distributions of sensorimotor events (Friston, 2008).

Predictive coding approaches are actually a step backward into internalism as is made clear in the next sub-section. In contrast, from a sensorimotor perspective, action-oriented representations are probably

[2] The only update in this respect on our previous reports is that we have recently implemented a model of torque compensation using spinal circuits in a multi-joint arm system that do not use internal models of any kind. See Buhrmann and Di Paolo (2014b).

the most philosophically sophisticated version of representationalism. They have nevertheless been subject to strong criticism. In particular, according to Dreyfus (2007) the action-oriented representational approach "has to face the frame problem head on" (p 358). In other words, how is it that action-oriented representations are able to be contextual, i.e., to track the context of relevance established by the current lived situation (states of the environment, motivations, intentions, states of the body, demands from others, etc.) in order to select the most adequate virtual action? This question remains unanswered for this and, to date, all representational approaches.

PPSMCs

Seth (2014) provides an interesting alternative, in that he builds on the predictive coding (or predictive processing, Seth treats the phrases as equivalent) model, but he also explicitly adds SMCs to this framework. He names the resulting approach the PPSMC model of experience.

Seth argues that when the hierarchical generative models (HGM) of the predictive coding account model counterfactual actions, then this accounts for the subjective reality of certain experiences. He takes as his paradigm case the distinction between synesthetic and normal experiences. His claim is that in normal experiences, a rich set of counterfactuals are modelled in the HGM (what would happen, if the subject interacted with the world in many different ways) and that it is the existence of this rich set of counterfactuals which accounts for the subjective reality of the experience.

Seth's key case of application of his approach is to synesthetic experience. He argues that the lack of subjective reality of synesthetic experiences (the subjects are aware of the synesthetic 'qualia', but don't mistake them for real items in the external world) can be explained by the lack of any internal modelling of the rich possibilities for counterfactual interaction.

Seth claims that one key advantage of his model is that it describes mastery (in the sense of practical understanding) of SMCs, something which O'Regan and Noë often talk about, but for which they do not offer any specific model. That is, O'Regan and Noë believe that experience consists in mastery of SMCs: practical sensitivity to possible SMCs in the current situation. Seth fleshes out this mastery in terms of sensitivity to counterfactual changes in SMCs (as does Beaton, 2013). But Seth argues that these counterfactual SMCs need to be explicitly modelled (in this case, within the HGMs of a predictive coding approach).

It is interesting to note that Metzinger (2014) praises Seth for having found a way to internalise SMCs. For Metzinger, it is a weakness of Noë's approach (in particular) that "interaction with the physical world is metaphysically necessary for presence" (Metzinger, 2014).

Froese (2014) and Di Paolo (2014) have exactly the opposite response. Froese believes that Seth's account could be sufficient to explain objecthood (the sense that the tomato has a hidden side, for instance), but that it cannot be sufficient to account for subjective reality. Froese follows Fish (2009) and Beaton (2013) in suggesting that subjective reality really does correspond to the involvement of the world itself in structuring experience. More specifically, but reaching the same conclusion, Di Paolo (2014) points out the inconsistencies and *ad-hoc* implicit assumptions regarding Seth's proposal. Seth seems to suppose that there is no interference between how probabilities are encoded in hierarchical generative models for actual processed stimuli and for counterfactuals leading him to suggest that lack of veridicality in synaesthetic experience is due to the independence of these codings and the fact that counterfactual coding could accumulate errors. It is interesting to link veridicality to counterfactual knowledge (very much in line with the world-involving proposal we will discuss later). But it is risky to assume these counterfactual encodings are unaffected by real sensory input and their relevant error-correcting mechanisms. It is in fact much more likely the opposite that is the case, i.e., that there is a strong dependence of any counterfactual knowledge on actual past sensory experience, for how could this counterfactual knowledge otherwise be constructed so as to be of any use for the agent? In other words, in Seth's proposal, counterfactual encodings must depend on actual coupling with the world. Hence the supposed internalization of SMCs fails and what Metzinger's sees as a weakness of the SMC theory is the inevitable, and correct, conclusion: the contribution of the world must be included in a correct account of sensory presence.

Below, we will follow up on these ideas. We will give a dynamical systems account which shows that mastery of (i.e., practical sensitivity to) SMCs is possible, without an internal model of the SMCs. And we will recap the philosophical account of Beaton (2013) (developed within the project) which argues that certain phenomenological features of experience (such as richness, and directness) can only be accounted for by considering the direct involvement of the world itself in experience.

2.2. Some approaches investigated in the project

All of the above approaches try to show that action and perception are linked through some intermediate sensitivity to the results of possible actions. In the accounts above, it is assumed that this requires some explicit model of whatever one is sensitive to. For instance, if one is to be sensitive to the way one's arm will move (in real time, faster than neural feedback would allow) it is assumed that one will need an internal model of the mechanisms of one's arm. Similarly, if one is to be sensitive to how SMCs change as one interacts with objects in the world, it is assumed that one needs a specific model of those objects, and of the resultant SMCs.

However, there has been a strong sense in the literature that SMCs may somehow be associated with non- or anti-representational approaches. The question is thus whether this association is a problem for the SMCs approach to virtual action, or whether a non-representational conception is after all viable. In the sections below, we firstly review several project architectures which take a strong step in the anti-representational direction, by showing how successful interaction with objects and the world can arise without modelling the world. The system can master the structure of the world, simply by mastering the structure of the relevant SMCs.

The UPF DAC-based approach

In a biologically inspired architecture called Distributed-Adaptive Control (DAC) (Sánchez-Fibla et al., 2011; Duff et al., 2011; Sánchez-Fibla et al., 2013), SMCs are conceived as *if-then* rules that match the occurrence of a discrete perceptual state to a discrete action to be taken. Such rules are represented by bi-modal memory units, of which at least one, but generally many, are assumed to exist for each possible perception-action tuple (though versions acquiring the memory only from actually enacted sensorimotor tuples can help deal with problems of combinatorial explosion). These bimodal memory units compete probabilistically for control of the motor output based on their activation value, which is primarily driven by each unit's sensory input but also modulated by reward. Thus perception-action rules that mostly received reward in the past are more likely to be selected for access to motor control.

It is to be noted that this architecture is explicitly thought of as representational. It is argued that the contextual layer of the DAC architecture (i.e. as just described) will model the presumed biological role of the prefrontal cortex: that of storing "representations of sensory states, actions and their combinations" (p.290). The aim of Duff et al. (2011) is to show that SMCs can be learnt in a biologically plausible (or, at least, biologically inspired) architecture. It should be noted that task-relevant inputs and actions are predefined, in the relatively abstract tasks used in that paper. On the other hand, other project architectures attempt to work out these task-relevant inputs and actions from lower-level primitives, as we will see below. However, the Duff et al. architecture certainly does successfully prove that the SMCs structure of various tasks can be learnt during interaction with the environment. Extensions to the approach show that object-related affordances can also be successfully learnt (see, e.g., Sánchez-Fibla et al., 2011). Once again, though, the background assumption of the approach is thoroughly representational (e.g. "we focus on ... the process of acquiring internal representations of object affordances", p.1).

It is interesting that Duff et al. note that: "Given the finite size of the genetic prespecification of biological entities it is reasonable to assume that knowledge needs to be acquired through interaction with the environment" (p.290). Indeed, recent research suggests that the number of genes in the human genome is so surprisingly low that it is likely that the genome itself does not fully code for the human body plan. That is, there is no full specification of the human body in our genes - our body form arises through an interactional process, in which the role of the environment is crucial. We can extend this point to suggest (contrary to the representational approach used in the DAC architecture) that goal-directed action itself may not be (fully) coded in the brain, but rather may also be a fundamentally interactional process - with the body and environment playing crucial roles. It may be that no (full) internal coding of the SMCs which the agent masters is required.

It is this possibility which we examine in this report. Next, we will discuss another project architecture which takes interesting steps in this direction.

The UKE Markov model of SMCs

In one of the key approaches developed within the eSMCs project (Maye and Engel, 2011; 2012a; 2012b; 2013), SMCs are learnt in a probabilistic manner. The principle is that, for any given action, a particular sensory result is more or less likely. If we take recent action-outcome history into account then we can make an even more refined estimate as to which sensory outcome is likely, for which current action.

In the Maye-Engel model, discrete-time h^{th} order Markov models are learnt, where h is the history length which the robot will consider. The robot generates actions and records the sensory consequences. It turns out that for a simple real (or simulated), embodied robot only a fairly small percentage of the action-outcome (*ao*) space is encountered. Thus, the robot is soon able to estimate the probabilities of action-outcome pairs, based on their observed frequencies so far. The robot also has an evaluation function. In the existing examples, appropriate evaluation functions have been handdesigned to achieve the required tasks, e.g. the robot may be 'punished' for tumbling (falling over), or for hitting anything with its bumper. Given this setup (essentially, a profit-loss function and a probability distribution) it is possible to choose a next action which maximizes the expected reward. Depending on the history-length used, this probabilistic evaluation may only depend on the previous timestep, or it may depend on a history several steps in length.

This approach has been employed in multiple successful robotic implementations within the project. To give a simple example, a simple simulated robot can easily be trained to learn the difference between a 'spike' and a 'block' (narrow and broad objects), and to track one but not the other (Maye and Engel, 2011). In another example, in both simulated and real robot platforms, it was possible to train the robot to move within its arena, whilst avoiding the arena walls (Maye and Engel, 2012a). Moving to more complex examples, in a both a real and a simulated Puppy robot (i.e. using four legged locomotion), the robot was able to learn which of several possible gait patterns was most stable on various different surface materials, and to adopt the gait (or pattern of gaits) which gave it the greatest overall stability (Hoffman et al., 2012).

It is important to note that, in all these cases, the environmental parameter of interest (e.g., what the supporting surface was; whether the robot was near or far from a wall) could not (even in principle) be read off directly from any sensor, but rather was implicit in the patterns of sensorimotor interaction over time.

From the perspective of the current discussion, then, this approach looks like an interesting intermediate case. One the one hand, the nature of the supporting surface (say, or the distance from the wall) is not explicitly modelled. Thus, it would be accurate to say that the robot's competence in these tasks is truly embodied – the robot can only continue to be competent when it is in an environment which affords the same types of sensorimotor contingencies as those which it has learnt. Furthermore, this competence is not entirely internal: successful completion of an action trajectory depends on the world, which continues to give its feedback and continues to affect the sensors in (at least broadly) the ways which the robot has learnt.

On the other hand, this approach does use explicit, internal models of action-commands and their corresponding sensory-outcomes. Thus, at a more fine grained level of detail, this robot masters SMCs by explicitly modelling SMCs. This is interesting. This Markov control model can master object-shapes, or gait-control, without explicitly representing these object shapes, or gait patterns. Yet, it seems, it can't master SMCs, without an explicit, internal representation of the SMCs themselves (at least at the fine level of detail).

This leads back to one of the central questions which we aim to address in this report. Is it possible to master SMCs, without any full, internal model of any aspect of these SMCs? To put it another way, can we really understand how to build a control architecture which involves the world (or, at least, can when required involve the world) 'all the way down'? We believe, and will argue, that the answer to these questions is affirmative. As second phrasing of our question, just above, might suggest: we believe that one can only understand how an affirmative answer might be possible, by taking into account the relevance of virtual actions to a truly embodied control system.

This is the approach which we will now examine. Firstly, we will briefly presenting some relevant, antecedent work. Then we will go on to introduce and discuss new developments.

3. The dynamical approach to virtual actions

3.1. Existing dynamical approaches

The above project-based approaches show that it is possible to master interaction with objects, without explicit modelling of the relevant objects. Instead, the creature's historical SMC interactions are explicitly modelled, and then an algorithm for selecting the highest value action is applied. The implicit sensitivity to objects emerges from an explicit model of the structure of the SMCs which the agent has encountered. But one might wonder, if it is possible to be sensitive to objects without modelling objects, is it possible to be sensitive to SMCs without modelling SMCs? Here, we examine some ecological and dynamical systems attempts at answering this question.

Virtual actions, which we have defined as actions not fully actualised but that in some way bear consequences for the agent and possibly the world, fall within the large domain of cognitive phenomena demanding, according to some, representational explanations. In the words of Clark and Toribio (1994), they are "representation-hungry" problems, i.e. problems that deal with absent elements (not directly perceivable or realised) or abstract ones (including unspecified general goals). The idea that the class of such representation-hungry problems is defined univocally or that, given a particular cognitive phenomenon, it is a straightforward matter to decide whether it is representation-hungry or not has come under considerable scrutiny and criticism, particular from dynamical systems, enactive and ecological approaches (e.g., Chemero, 2009, Di Paolo et. al 2010, Cuffari et al. 2014, Degenaar and Myin, 2014).

To cite a couple of examples, Izquierdo and Di Paolo (2005) have provided a simple demonstration of how a type 2 classification problem (one in which the instantaneous stimulus is ambiguous with respect to the correct moment to moment decision about the course of action to take, Clark and Thornton, 1997) can actually be solved by a robotic agent using only a stateless machine: thus incapable of handling any complex representation about the history of stimulation which is necessary to extract and internally store the relevant input information to solve the classification task. The key aspect to the non-representational solution is that the agent must be able to engage in a continuous dynamical coupling with the environment, and it is in this dynamical process that the ambiguities of the sensory input are resolved. This follows on from similar, and remarkable, demonstrations of simple embodied agents performing supposedly representation hungry tasks such as delay, history contingent choices, learning and conditional classification using the minimal cognition paradigm championed by Randall Beer and colleagues (e.g., Beer, 2003, Izquierdo et al. 2008, Williams et al. 2008).

While Clark and Toribio (1994) (see also Clark 1997) have in mind what we could roughly describe as higher cognitive mental function, when talking about representation-hungry problems, their proposed category would seem to include, for some, any kind of motor control problem for which direct sensorimotor feedback would seem insufficient (see Deliverables D1.1 and D1.2 for discussion on the use of forward and inverse models in motor control). For instance, performing smooth arm movements (or smooth movements in any other multi-joint system) involves compensating for interaction torques which emerge on a given joint due to the movement around another joint. Controlling for interaction torques would seem to be a representation-hungry problem because sensory feedback can be ambiguous (i.e., relevant information would seem to be absent). However, within the eSMCs project, Buhrmann & Di Paolo (2014b) have demonstrated in a model implementing simple spinal circuit equilibrium-point control that the dynamics of the musculoskeletal system is able to compensate by itself for such interaction torques.

In short, what is or is not a representation-hungry task is anything but a trivial or even a decidable matter. We conclude that the category is best dropped altogether (see also Degenaar & Myin, 2014).

How does this affect the notion of virtual action? Firstly by introducing the following clarification: The fact that virtual actions involve non-actualised aspects of actions (e.g., movements not taken) does not *a priori* imply that non-representational explanations must be discarded by default. So far this seems simply a weak clarification, but it gains weight in the light of the fact that this is precisely what most cognitive scientists and neuroscientists do: they assume that there is no other way to solve cognitive tasks other than through representations. Such assumption is made by default, often offering zero evidence and without the aid of an argument.

It remains however to be seen whether non-representational approaches can provide competitive alternative explanations for virtual action.

Let us consider a task that we may describe in terms of virtual actions. Such as task was studied empirically by van Rooij et al. (2002). They provided participants with sticks of different lengths (either increasing or decreasing in length in sequence) and ask them to imagine whether they could move a distant object with the stick they've been given. In other words, the participants are asked to predict the outcome of a virtual action, one they do not actualise.

The authors proposed a simple dynamical system model to describe the participant's behaviour. According to this model (first proposed by Tuller et al. 1994 for studying a speech categorization task), a collective variable *V(x)* is used to describe the evolution of the 2 choice option given to the participant at each moment of the ascending or descending sequence of sticks: To answer 'yes' or 'no' to the question: do they think they can move the distal object?

V(x) describes the overall emerging behaviour of the system (participant + stick + object) according to single parameter *k* according to

$$V(x) = k\,x - \tfrac{1}{2}\,x^2 + \tfrac{1}{4}\,x^4.$$

The effect of *k* is shown in Figure 1.

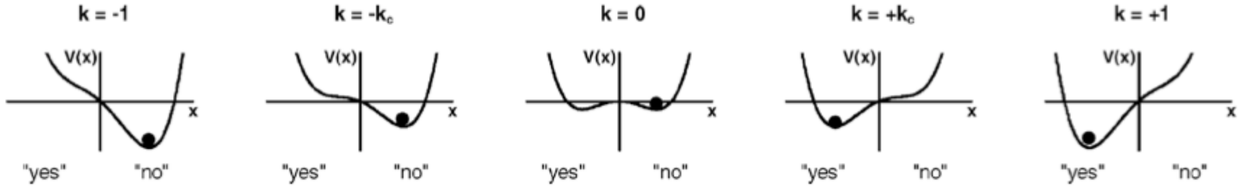

Figure 1: potential function describing the participant's response in van Rooij et al.'s (2002) virtual action experiment (copied from the original).

The control parameter k is made to vary according to: the length of the stick at a given trial, the length of the stick at the previous trial, and the participant's Yes or No responses at previous trials.

Based on this dynamical model, the authors made four predictions. Firstly that there would be an assimilative bias leading participants to tend to repeat the response of the previous trial. Second, there will be an inverse relation between the length of the stick on the previous trial and the probability of answering yes in the current trial, a contrastive effect. Third, the model predicts hysteresis and middle range stick lengths may produce a positive or negative answer depending on whether they are given to the participant in an ascending or descending order of presentation. Finally, the number of ambiguous length leading to the probability of providing different answers in different occasions depends on the size of the multi-stable region in intermediate values of k (see figure 1).

These predictions were all verified empirically. The authors discuss that while it is possible to imagine representational accounts that would explain some of the predictions, these accounts would tend to contradict other predictions at the same time (e.g., hysteresis and enhanced contrast) leading to overly complex models, when a simple dynamical system describes the emerging pattern of behaviour more accurately.

While such experiments provide a strong indication of the importance of studying virtual actions in a situated, world-involving model comprising agent, world and history, as opposed to in-the-head computations, these experiments still lie mostly on the descriptive side, and raise the question of what kind of explanatory mechanisms might be involved in producing virtual actions.

We return to this question at a later section, but to finish this exposition of non-representational ideas, we will simply mention that the notion of anticipation has been the object of much study in ecological, dynamical approaches, and that such a notion conceives of the historically situated agent as "poised" in a metastable state which can quite rapidly crystallize into the right cognitive performance. The stimulus acts as a trigger and not so much as bearing information that needs to be processed and decoded. Evidence of this idea is provided by cases of so-called ultrafast cognition (Wallot and van Orden, 2012). The authors review evidence of fast cognitive performance (typically in go/no-go tasks) that occurs reliably in times which are faster or of the same order as the estimated physiological limits of information processing (see also Thorpe, 2002).

Evidence of ultrafast cognition supports the notion that cognitive systems can be *strongly anticipative*. According to Stepp and Turvey, (2010) strong anticipation arises from the (ecologically situated) system itself via lawful regularities embedded in the system's ordinary mode of functioning and in coupling with the world. Weak anticipation, by contrast, arises from a model of the system via internal simulation. The distinction is introduced in formal terms by Dubois (2003). A strongly anticipative system is already attuned to its situation by the integration of feedback, delays and synergies into its overall dynamics. Before movement starts, the system is already in some form of attractor that will invariable lead to the appropriate, context-sensitive behaviour. Perturbations, like pulling on the chin while you're trying to articulate a phoneme (Kelso et al. 1984), simply result in taking a different transient towards that same attractor (equifinality with respect to the outcome, rather than the implementation of movement). In other words, the existence of the attractor means that the system already anticipates possible perturbations by allowing for different instant changes in its transients. And the setting up of the attractor is a sort of feed-forward action, hence by the time the movement actually starts, nothing more needs to be "computed", the system just relaxes towards the fixed-point (where the system could of course be the agent-environment coupled system).

This relaxation can indeed be context and history sensitive as demonstrated in Buhrmann and Di Paolo, (2014a), where an arm model with a touch sensor at the tip is able to follow a surface in front of it at various orientations and positions and "continue" moving along the same plane even when the sensor is turned off (as if predicting where the remainder of the surface "should be"). These trajectories, however, all lead as different transients towards the same attractor (Figure 2).

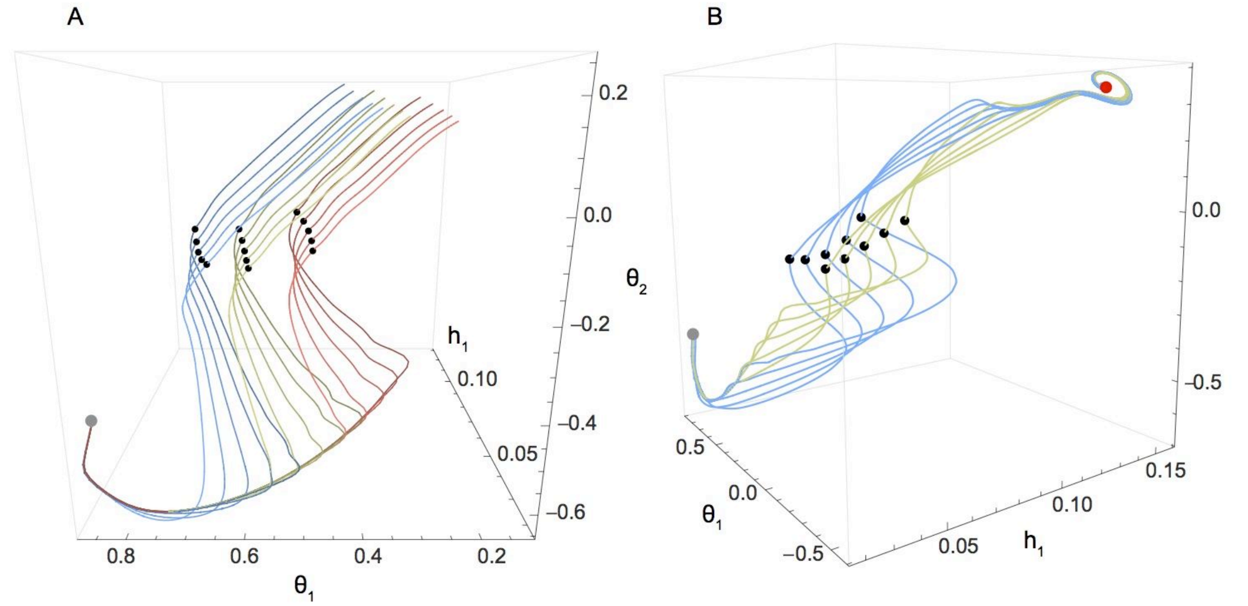

Figure 2: Trajectories of arm model in Buhrmann & Di Paolo, (2014a). Three-dimensional projection of the agent's neural state onto the two proprioceptive input neurons ($\theta_{1,2}$) and a hidden neuron (h_1). A: trajectories for three different surface positions (blue, green and red) and five different orientations each (shading of colour changes with surface angle). Black markers indicate the point of touch sensor deactivation and a grey marker the common initial state. B: Trajectories for two surface positions over a period of 50 s. A red marker indicates the common attractor position.

3.2. Interlude

Let us recapitulate some of the main points of the above mentioned approaches to virtual actions. As we have suggested, the notion of an action not taken but that somehow has a cognitive and perhaps even a worldly effect was originally associated in traditional AI with research on action planning. All of the relevant GOFAI work using symbolic representations and logic programming in action planning classifies as strongly representational. At the other end of the spectrum, some of the ecological and dynamical approaches, though still in need of further development (see next section), claim that there is no need for a representational duplication of what the skillful agent is already able to perform in coupling with the world because this performance is already oriented toward action possibilities, even if these are not taken.

This distinction between representational and non-representational approaches can be further refined in two related ways. One refinement looks at the role of embodied action in shaping representations in form and content. Accordingly, traditional representational approaches base their planning on an objective, geo-centric model that establishes relations between represented external objects based on rules describing their interaction. These models can be causal, folk-physical, probabilistic, etc., but they are characterised by an independence of representational form and content on the details of the agent's sensorimotor embodiment. Other representational approaches are, as we have seen, action-oriented. They make use of ego-centric, body-formatted representations of action possibilities. They could be more abstract (in terms of schemas, moves, etc.) or more concrete (in terms of sensorimotor contingencies, sensorimotor predictions, etc.). Or they could involve both possibilities in an hierarchical modelling structure.

Another dimension of refinement looks at the role played by the world. Traditional approaches limit the role of the world to that of an informational input. The burden of cognitive processing lies exclusively with the "agent's head". We can call them "in-the-head" approaches. Other perspectives allow the world to play more than a simple informational role, i.e., the dynamics of body and world are constitutive of the cognitive processes involved in virtual actions; they cannot be reduced to providing an input to the agent's cognitive processes. We call these approaches "world-involving". Notice that world-involving does not mean constituted *exclusively* by the world. World-involving proposals also rely on processes internal to the agent as well. World-involving, in other words, means not-*just*-in-the-head.

There isn't an obvious simple mapping between representational (whether abstract or action-oriented) and non-representational approaches and in-the-head and world-involving approaches. Wheeler's (2005) action-oriented approach involves both representations (in terms of action possibilities) and the dynamics of the world (in terms of capitalizing on body-environment dynamical couplings). It is arguably world-involving, but places the greater part of the cognitive load of planning in-the-head. Seth's (2014) SMC-PP approach, in contrast, while based on probabilistic models of sensorimotor expectations (i.e., shaped by experience in the world) is not world-involving, since the world figures merely as input to the lower sensorimotor levels of the hierarchical models. It is an in-the-head approach.

Finally, we note that while a world-involving perspective may or may not be representational (they world may play exclusively the role of informational input, or it may play other constitutive roles, including or not that of informational input), any non-representational approach must *per force* be world-involving. The reason is that in any non-representational approach, cognitive performance must be constituted in terms that are different from those of inputs to information processing (otherwise it would involve some kind of representation). This means any such approach always involves some non-informational aspects of the coupling with the world.

3.3. Developing an enactive approach to virtual actions

We have already suggested that attempts such as Seth's (2014) at combining SMCT with predictive coding (or similar emulation or internal models approaches) would possibly offer an account of virtual actions which is very much in-the-head and not world-involving (in the strict meaning given above to these terms). In many senses what these moves offer is a sensorimotor "update" to more traditional representational theories. This in itself may be a good idea, of course. But we think (and would like to explore this here) that it may be possible to pursue a more radical alternative, one that that we imagine more compatible with the spirit of the sensorimotor approach as expressed repeatedly in the primary literature (O'Regan & Noë, 2001, Noë, 2004), an account of virtual actions that is *world-involving.*

For this, the different dynamical and ecological proposals briefly reviewed above offer a good starting point. But before we proceed, let us look closely as the problems that such a proposal should be able to face.

If we were to offer a world-involving (possibly non-representational) account of virtual actions using dynamical, ecological and sensorimotor theories, are we not faced at the very start with a contradiction in terms? After all, we have defined virtual actions as actions not taken, not actualised, but that somehow leave the agent altered and allow it potentially to master more complex forms of sensorimotor activity (like the planning of complex sequences of action before their execution). But how can these virtual actions not be actualised and the account offered for them be at the same time world-involving? A world-involving account implies that our explanations involve the world in a non-trivial way (i.e., not merely as a source of informational input but as causally enabling and possibly constitutive of cognitive phenomena). Normally, in world-involving sensorimotor accounts, the world is involved non-trivially, but precisely because the behaviour accounted for is actualised, i.e., it is indeed happening *in coupling* with the world. This is the explanatory "entry point" for worldly factors to play a non-trivial role.

A world-involving account of virtual actions would not have it so easy. Indeed, it would seem as if by definition, any account of virtual action is placed within the ill-defined category of representation-hungry problems discussed earlier. Since virtual actions do not take place in actuality, in what ways can we say that the world is involved in them taking place at all?

Here we are simply making explicit the default reasoning that motivates all representational, in-the-head approaches.

The first point of clarification that should be made is this: *the coupling between agent and world never disappears, and the sense-making that it supports is not restricted only to sensitivities in the here-and-now.*

Critics of enactivism worry that enactive accounts of sense-making can merely handle low-level and immediately present or concrete phenomena. The typical move (already discussed above) is then to separate cognition into dynamic, embodied, sensorimotor accounts for low-level forms of "online" activity (grasping objects, performing some basic perceptual discrimination), and representational accounts for anything more complex than that, i.e., so called "off-line" activity. The enactive perspective sets aside such a distinction between online and off-line as meaningless (Cuffari et al., 2014; De Jaegher & Di Paolo, 2013). The point is expressed in terms relevant to the current discussion in Cuffari et al.'s enactive account of language. According to them, critics of enactive accounts of language

> " … presuppose a scenario in which an individual cognizer is not necessarily in constant co-constitutive relation with her environment. Yet these imagined cases are built on an erroneous conflation of the operational conditions of cognitive processes with the *meaning* achieved by a cognizer thanks to those processes, and on an erroneously reductive understanding of what it means for something to be 'right there'. It is also crucial to note that the online/offline dichotomy can only be articulated according to the logic of mental representations .... The end-game of such an account is unavoidably a return to a disembodied, computational-functionalist model of cognition. In contrast, the point of our story is to offer an interpretation of 'living system' and 'niche' that locates meaning, including what may appear to us uncritically as 'internal', 'detached', 'offline' or 'abstract' meaning, in that interaction." (p. 6).

They continue to state that:

> "... though we may describe our cognitive abilities as transporting us 'beyond' the present moment, it is misguided to attempt to locate the cognitive abilities *themselves* somewhere beyond the present moment. [The critics'] fallacy of misplaced concreteness is to associate meaning with coupling. This makes them wonder how we could ever mean what we are not coupled to. Meaning, in contrast, is the relational activity of sense-making, which holistically involves the autonomous agent's adaptive modulation of its own dynamical tendencies and its coupling with the world (Di Paolo 2005, 2009; Thompson 2007). *Never* in any of the descriptions of sense-making in the enactive literature has meaning been equated with coupling. The coupling between variables and parameters between agent and environment is only one element in the sense-making process, certainly not the bearer of any meaning by itself. Human sense-making involves a range of sensitivities, importantly including linguistic sensitivities, and so the human niche, that which we couple to, is a linguistically mediated and layered or 'enlanguaged' world." (ibid.)

In other words, the tension between world-involvement and the virtuality of non-actualised actions is equated to the tension between online and offline forms of cognition. But this tension dissolves once we stop trying to explain cognitive content as being carried in vehicles. Traditional approaches think along these lines: If not carried within internal representational vehicles, meaningful content must be "carried" somewhere else, perhaps in the coupling with the world. But the latter being apparently insufficient to carry content related to absent aspects of the world or non-actualised aspects of the agent, it is therefore to be discarded as an explanation of virtual action. The whole argument however becomes unnecessary once that we stop thinking of cognition as the management of content that must be carried from one place to another (see also the relevant discussion in Hutto and Myin, 2013).

Having established that attempting a non-representational and world-involving sensorimotor account of virtual actions is not doomed from the start, we can proceed with sketching some of the main points of this account. This will by no means be a complete proposal. It should be considered as work in progress.

The first point to consider is that from an enactive perspective any form of action and any form of perception is a form of sense-making. To repeat what has been expressed elsewhere (Deliverables D1.1, D1.2), sense-making is adaptive regulative activity of a self-sustaining autonomous and precarious organism such that through this activity the system is sensitive to trends in its states and relation with the world that would break its boundary of viability and in turn responds by attempting to regulate these trends such that such a breakdown is avoided. Let us expand on this point.

The autonomous organism is capable of evaluating a given interaction with the world as beneficial or not for its own continuation. In other words, the maintenance of systemic identity is a source of intrinsic normativity: eat food, avoid poison, etc. The problem is that the formal condition of closure of a living system as presented by autopoietic theory (the operationally closed network of constitutive processes) is in itself binary, not graded. The network of process relations is either closed or not, and as a consequence it cannot serve directly as a source of graded values (e.g., this is better or worse than that) to filter through encounters with the world. Some form of mediation is missing in this picture, something that can link actual states with the tendencies and capacities involved in evaluating the consequences for the system's own viability. Sense-making requires not only an autonomous organization but also an adaptive one (Di Paolo, 2005). Adaptivity is the capability of an autonomous system to respond to tendencies in the trajectories of its states and its relations to the world, such that when these tendencies approach the boundary of its own viability the system modulates its coupling with the world in a way that tends to avert the crossing of this boundary. Adaptive responses thus permit the regulation of states before the breakdown of closure (death) has occurred. Such responses can succeed or fail, which is crucial as this makes adaptivity a graded property that naturalizes sense-making in the system's organization.

Relevant to our discussion is the fact that adaptivity works on the *virtual field* that surrounds the current dynamical configuration of the agent–world system. In two senses: first, the agent must be responsive to whether or not tendencies approach the boundary of viability — if the crossing of this boundary is actualized, it is simply too late; second, it must make use of its capacities to modulate the constraints of the coupling with the world by introducing changes that alter the virtual field around the current states (modifying the direction of the negative tendencies).

What exactly do we mean by a *virtual field*? The notion of the virtual in the current context follows from the metaphysical tradition of granting a real ontological status to capacities and tendencies, especially those in the neighbourhood of current actual states. This tradition includes Aristotle's notion of potentiality, Spinoza's concept of affect, and related (but not identical) usages of the term virtual by

Bergson and Deleuze. For our purposes, we additionally use the rich conceptual resources of dynamical systems theory to speak of trajectories, traces, and tendencies. In a few words, in this tradition, the virtual is that which is real but not actual, such as the glass's capacity to hold liquids or the tendency of liquid water to become solid at freezing temperatures. Such capacities and tendencies need not be currently actualized (the glass is empty and in principle could remain empty forever but its capacity is still real).

Notably, these capacities and tendencies are always *relational*, unlike the properties of an object that belong to it in all contexts (e.g. water is composed by H_2O molecules). Hence, virtual capacities and tendencies can be potentially infinite in number (the glass can be used as a paperweight or a doorstop, water also has a tendency to slow down high-energy neutrons inside a nuclear reactor). This doesn't mean that we can manipulate circumstances so that any object could acquire *any* arbitrary capacity or tendency (the glass hasn't got the tendency to slow down fast neutrons, and liquid water makes a rather ineffective paperweight and doorstop). Actualization of virtual capacities and tendencies occurs always as an *event*, or a doing, or an act. These actualizations are situated in history, have duration, and so forth.

By speaking of *fields* we want to focus on the structured situatedness of virtual tendencies and capacities, i.e. not just any non-actualized counterfactual (though see discussion in the next section), but those that neighbour a current state of affairs. This adjacent structured field can be studied by examining the dynamical landscape around a current trajectory (this landscape in itself need not be fixed in time). Such fields have been investigated in many simulation models of embodied cognition to different degrees of explicitness. It is common in such models to artificially travel back and forth along a recorded behavioural trajectory and manipulate circumstances and re-run the model to ask 'what-if' questions in order to systematically map a virtual field (e.g. Beer, 2003; Iizuka and Di Paolo, 2007) or to attempt to produce exhaustive maps of sensorimotor possibilities (Buhrmann, et al., 2013).

We come to our first intermediate conclusion: in a strict sense, all of sense-making, *a fortiori* all action and perception, involves constitutively an element of virtuality.

In other words, sense-making, which even in its most basic forms implies adaptivity, always occurs in a 'thick' here-and-now. By this we mean that, given the current situation, not only the actualized states matter to the sense-maker but also the virtual traces and tendencies that surround these states, whether they become actualized or simply modified but not necessarily actualized in the course of events.

From a dynamical systems vantage point, this is no surprise and here we can see the affinity with the dynamical models discussed earlier. Let us consider three widespread characteristics that become obvious when agents are considered in dynamical terms. 1) The current state of a system reflects a history of changes that the system has undergone over time. In this way, past events are brought to bear on the current situation and so the accumulation of experience allows the agent to discriminate between different contexts when exposed to identical sensory perturbations. 2) The behaviour of the agent as a dynamical system depends on its limit sets, i.e. the macro-configuration of states that need not be actually ever visited by the system. Since through continued coupling with the environment the agent is able to reach different areas of state space, from different initial conditions the agent may then follow different tendencies as determined by the attracting and repelling sets. Experience can therefore tune those limit sets globally such that the agent's movement through state space corresponds to the desired response that a given situation solicits. 3) Future tendencies depend on how the agent's dynamics moves towards the limit set in whose basin it finds itself; asymptotic states need not ever be reached (actualized) or even known, but they nevertheless exert an influence in these tendencies.

One by one, these properties bear a correspondence with three key aspects of Merleau-Ponty's notion of motor intentionality (Buhrmann, et al., 2013; Buhrmann and Di Paolo, 2014a). 1) The accumulation of experience serves to discriminate, with increasing specificity, situations that solicit a particular response (separation of traces). 2) Experience also allows a person to incrementally refine her dispositions to respond to these solicitations (tuning of tendencies by altering limit sets). 3) A response to a situation takes the form of movement toward the completion of a Gestalt ('maximum grip') or equilibrium towards which the body tends without the need to mentally represent this optimum, as when finding the right distance to admire a painting (dynamics in a basin of attraction).

In the enactive analysis, making sense of a situation — i.e. evaluating its relevance to the sustaining of a precarious identity and responding accordingly — requires by definition both *sensitivity* to and *modulation* of the virtual field of possibilities. The sensitivity and capacity for modulation are present in all forms of life with large variations in complexity and refinement. For the average person, it is not the same to walk a straight course stepping along the curb at street level or stepping along the border of the terrace of a tall building, even if the sensorimotor trajectories that should be actualized are nearly identical. For the average pigeon the two situations are similar. This is because both are differently sensitive to risks and have different capacities to respond to them.

We see then that normal action, from an enactive perspective, involves both sensitivity to the virtual adjacent field of possibilities and a way of manipulating this virtual field (for instance by reducing risk).

Our proposal is that within this world-involving account then, virtual actions involve the skill of extending such sensitivities further in time and recursively involve in such sensitivities the virtual consequences of actions not yet taken.

This skill, for which we acknowledge that a more detailed operational account is not yet available, involves no new fundamental principle over and above the concept of sense-making. It does require, yes, more complexity. Consider how we have discussed the notion of mastery of SMCs. This can be recast in terms of sensitivities to the virtual field conditioned on possible actions that are not immediately taken. I see a mug in front of me and my perception involves constitutively my mastery of the virtual possibilities afforded by this mug and my body in interaction. I may move my head slightly or reach to grab the mug and my mastery of SMCs informs my perception of the mug by the sensorimotor characteristics of these actions even if I don't take them. To perceive in the sensorimotor approach is already to master the consequences of virtual actions.

What is required to account for more complex forms of virtual actions is an account that furthers this mastery in time and into more complex forms for virtual sensitivities. The immediately present mug is like a "hub" for directly solicited virtual actions (move the head around it, grasp it, drink from it). A skillful sense-maker could regulate these immediate sensitivities to include virtual events further in time (grasp the mug now but then place it on top of another surface, near the desk at which I will sit while I write a report).

There is no fundamental difficulty here, only the complexity involved in skillfully extending and navigating the virtual possibilities afforded by the here-and-now. In such an account, we could even speculate on how to define "objects" from a sensorimotor perspective. They would be "hubs" of ramifying virtual actions which would conserve some invariant characteristics across all of these actions.

We note that this proposal is highly compatible (though not identical) with ecological approaches. Indeed the notion of affordance is a paradigmatic case of a real virtual relation between organism and world that does not need to be actualised in order to be perceived and acted upon. Indeed, a recent proposal by Brincker (2014) elaborates a very similar perspective using the notion of affordance navigation to give an account of complex cognitive phenomena such as the passing of false-belief tests.

4. Virtual action and first-person experience

In this section, we develop a final strand of the relevance of virtual actions to the philosophy of SMCs. The discussion here covers related ground to the issues above, of action, representation and normativity. But now we try to squarely face the problem of how the SMC approach should account for first-person experience. As we will see, several important, common themes emerge. This last section draws largely on work published in Beaton (2013), which was completed as part of the eSMCs project.

Neurophenomenology, as introduced by Varela (1996), has been influential in the development of the enactive approach (Gallagher & Zahavi, 2008). Neurophenomenology aims to seek a mutually guiding two way interaction between "disciplined first-person accounts" of experience, and neurobiology. In fact, the approach is not limited just to neurobiology. As Varela says:

> "'neuro' refers here to the entire array of scientific correlates which are relevant in cognitive science. But to speak of a neuro-psycho-evolutionary- phenomenology would be unduly cumbersome." (Varela, 1996 p.330).

We will examine this point in this section. We argue elsewhere (Beaton, 2013) that researchers in the enactivist tradition have typically assumed that the structure of experience corresponds most directly to the structure of neural dynamics, despite Varela's quote above. Indeed, the suggestion that the structure of experience corresponds to the structure of external interactions (such that experience itself involves the world itself) is typically considered controversial.

Here, we argue that the structure of experience instead corresponds to the structure of *actual and available actions* which the experiencing agent takes or might take. That is, in this section we present a philosophical account of experience in terms of action. At first this might appear to be no more than the standard SMC approach; but here the novel aspect of our treatment is that we are much more explicit about the indispensable role of counterfactual action for the successful development of a sensorimotor treatment of experience.

A note about terminology: in this section, we will use the terminology "counterfactual" (which is clarified below), since this corresponds to pre-existing usage both within the project (Beaton, 2013) and elsewhere (Seth, 2014). But the notion of 'counterfactual' as used here has much in common with the Deleuzian notion of "virtual" action: action available, but not taken; action which is on the periphery of – but directly relevant to – experience. As we will see shortly, 'counterfactual' action – in the sense used in this section – is also quite directly relevant to resolving some of the central issues of the "Virtual Actions" task of the eSMCs project, such as the apparent problem of locked-in syndrome (LIS) for a sensorimotor account.

In SMC theory, a favourite example of the relation between action and experience comes from the study of perspectival vision, and more generally of vision and action in a 3D space (see Noë, 2004 Ch.3). Thus, for instance, a nearby object subtends a greater angle at the eye, compared to a further object. This corresponds to a difference in the types of action required, in order to point to, fixate on, reach out to, etc., near vs. far objects. Similarly, objects project different outlines as we move around them, and they around us. For instance, a penny viewed at an angle looks elliptical. This corresponds, for instance, to the fact that the action needed to correctly point towards (or look at, or otherwise trace out) the outline of a penny viewed at an angle is the same as the action required to trace out the outline of an ellipse viewed face on. In this sense, there is an objective part of the claim that these two look the same. More generally, SMC theory would claim that experience of solid, 3D shape is mastery of the actions required to trace out, grasp, fixate on, etc. the changing outlines, 'from here', of the solid 3D objects in our world. (Mastery means, the ability to successfully use such actions in norm-involving plans or schemes to which the object shapes are relevant.)

Often, SMC theory seems slightly stuck at this point, as if all the one can say is that one is experiencing shape (say) when one is engaging in the actions required to interact with shape. If this is right, it would imply that when one is not engaging in any overt, physical action, one cannot be experiencing anything. There are two possible lines of response to this, on behalf of SMC theory. One is to broaden the notion of action, to include certain internal changes of state as action (Kyselo and Di Paolo, 2014). Given the relevance of internal change to maintenance of system autonomy and viability, this route certainly fits well with the enactive approach.

But it should be noted that if we include the relevance of counterfactual actions, then the SMC analysis of perspectival experience in terms of action can succeed, without extending the notion of action to include internal changes in state. Or perhaps it is better to say, internal changes of state are relevant to experience but – it can be argued – they are relevant because they have at least counterfactual upshots in the action the agent would take. So, we say an agent is experiencing (the shape of) an elephant, if we have reason to think that the agent would in fact act in certain elephant-shaped ways, if appropriately tested. This analysis of experience in terms of counterfactual action renders experience less internal and mysterious than is often supposed. What another person is experiencing is something which you can know, if you perform the right behavioural tests. Or, at least, which you counterfactually could have known, if you had performed the right behavioural tests at the time.

4.1. Action vs. movement

This approach might at first sound similar to behaviourism, but it is not the same. Firstly, standard behaviourism actually denied the possibility of the scientific study of internal experience; it stated however that this was not a lack as the relevant object of psychology was observable behaviour. The approach argued for here, on the other hand, certainly accepts that internal, subjective, first-person experience exists and that its study is both possible and relevant. In fact, we argue that experience can genuinely be studied from the third-person as well as the first-person, because first-person experience corresponds directly to the structure of actual *and available* action. Secondly, behaviourism argued that all of psychology, including experience itself (to the extent, arguably zero, that anything recognisable as experience is supposed to remain present in the behaviourist approach) can be re-expressed in terms of meaningless, mechanical movements. This is one standard meaning of the idea of *operationalising* an object of study, i.e. one defines what is to be studied in a purely mechanical way. The approach we discuss here is not operationalising experience in this way. We are defining experience in terms of *action*, not in terms of mere movement (which would be operationalised, effectively meaningless 'behaviour', in the behaviourists' sense of 'behaviour'). Instead, we are defining experience in terms of action which, by definition (in the sense in which we are using it), involves norms: it is done for a reason, it is not meaningless.

It remains controversial whether it is possible to operationalise (in the sense just discussed) norm-involving action. Is there a purely mechanical way to identify which parts of nature have norms? Or do you have to presuppose norms - or even to have your own norms - in order to correctly identify norms in other parts of nature? (Enactive approaches attempt to solve this problem by proposing an operational description of the relations between behaviour and the agent's autonomy as a self-sustaining identity: organic, cognitive, social, etc.) Whatever the answer to these questions, the approach here defines experience in terms of the structure of actual and potential *action* (movement done, or potentially done, for the agent's own reasons), not in terms of actual and potential meaningless movement, as was the case with behaviourism.

4.2. All of experience?

Another possible objection to the action-based approach to experience would be that it cannot capture all of experience. Thus – it could be argued – the action-based approach might successfully capture the structure of our experience as far as the 3D structure of shape goes, but surely it cannot capture our experience of colour or smell?

A brief summary of the answer to this question involves two parts. The first part is to note that theorists working on SMC theory have indeed made progress on the structure of the actions which need to be mastered in order to perceive colour (Philipona and O'Regan, 2006; 2008). What is required is an understanding of the way in which lights and surfaces interact. No one claims that we need a theoretical understanding of these issues, in order to experience colour. But an SMC theorist can claim that we need a practical mastery of these issues, in order to experience colour.

Is this enough to capture all of experience, even the qualia? Beaton (2009) argued that emotion and affect is also relevant. But emotion and affect can also (arguably) be captured in terms of actual and counterfactual action. For instance, the most basic physical manifestation of emotional valence corresponds to approach versus avoidance behaviour. More generally, there is a rich history to the idea that emotions are fundamentally physical, in their most basic expression (James, 1884/1968; Darwin, 1872/1998). See also Colombetti (2014).

Thus, even if we need to include emotions in our account of experience, it is at least arguable that we can still do so without going beyond norm-laden actual and counterfactual behaviour.

4.3. Locked-in syndrome

The case of locked-in syndrome (Laureys, 2005) is often considered an insuperable challenge for SMC theory, because these patients are experiencing, but they are not moving. Worse than that, they cannot move, so even the counterfactual action theory might seem to have insuperable problems, here.

But, we will argue, the counterfactual approach can overcome this objection, and preserve an action-based, sensorimotor account of experience, even for locked-in patients. For this to work, we have to be able to claim that LIS patients do still know how to physically interact with their world, and would still do so, if only certain counterfactuals (i.e., things which are not true, but might be) obtained.

In this sense, then, there are at least two ways in which an LIS patient would physically, actually, act - if only they could. Firstly, we could argue that such patients would act again, if given access to brain-computer interfaces. At the current level of sophistication, this might only amount to control of a mouse on a computer screen, but when LIS patients are given means of communication (if only through eye-blinks and alphabet boards) they can show that they are aware of much more, and can carry out much more complex interactions indirectly, e.g., by making requests of their carers. Secondly, we can note that the brain damage in LIS is often very localised (deep in the brain stem). We can suggest that, if only this localised damage were reversed (although this is well beyond modern surgical intervention), then the patient would once again literally take the actions which they still know how to take.

It might sound as if this defence of the action-based view of experience threatens to descend into triviality. Surely even a fully vegetative patient could successfully act again, if only their brain were replaced by a fully functioning new one or if only the extensive damage to their brain were somehow repaired or reversed. In response to this objection, we can note that in the case of LIS, but not in the case of the full-blown persistent vegetative state, it is the physical locked-in patient, without any counterfactual additions, who still knows how to act. How so? In the locked-in case only, what the patient wants to do (would do, if only they were enabled to) is up to them. More reductively, it is physically determined by their physical state, as it is now (i.e., as it is actually, not counterfactually). Thus the interesting parts of the relevant counterfactuals (i.e., the facts about what the patient would do, if only they could) are already determined by the actual state of the patient, even though they may only be realised counterfactually. This is why it is right to say that the locked-in patient who still thinks, still has beliefs and desires, still wants, and knows, how to act.

4.4. Phenomenological advantages of the approach: Richness and directness

Thus, it seems, the inclusion of virtual or counterfactual action can allow an action-based approach to deal with some challenges which are often supposed to be beyond it. But actually, we will argue, treating experience in terms of (actual and counterfactual) bodily interaction with the world has additional advantages, in terms of enabling us to find a match between the first-person phenomenology of experience and the world which can be studied by science. Once again, we are going to argue that the structure of actual and available action matches the structure of experience better than does the structure of brain dynamics. In fact, in this case, we are going to argue that the structure of actual and available embodied, world-involving action can match the phenomenology of experience, whereas the structure of purely internal brain dynamics cannot be a good match for how experience seems to be from the first person.

Directness

Firstly, experience seems to be an encounter with the world itself. As Merleau-Ponty puts it (and cf. the Heidegger quote given earlier):

> "When we come back to phenomena we find, as a basic layer of experience, … not sensations … , but the features, the layout of a landscape or a word" (Merleau-Ponty, 1962 p.25)

But nothing in brain dynamics can be an encounter with (things in) the world itself. So if experience corresponds to the brain dynamics, then experience cannot be the encounter with the world that it seems to be.

In this case, internalist views have a further philosophical problem. If experience corresponds directly only to brain dynamics, then it seems that all we can ever be sure about are those brain dynamics themselves (under their experiential guise). That is, we would not really be in contact with the world, we could only be directly sure of our experience. Many, since Descartes, have thought that this sceptical conclusion is compulsory. But it is not the only possible view. Not just within phenomenology but even within modern Western analytic philosophy, alternative views, which claim that experience literally involves the world, have been put forward by very well-known authors (for example, McDowell 1996; Martin 2006).

If this alternative view is right, if the world really is a constitutive part of our experience (and not just a constitutive part of the causal story leading to our experience), then we are no longer left taking the world on trust. When everything occurs normally, and when we take ourselves to know that a table is right there, in front of us, on the basis of our experience, we really do have knowledge that is based on the table itself. Our perceptual experience (and hence our knowledge) constitutively involves the table, just as it seems to.

Richness

A second, related, aspect of experience for which brain dynamics alone can never be a successful match is its richness. Experience presents the world to me as transcending what I know. Experience outruns us, surprises us. There is always more to find. To take an almost over-simplified example, consider the visual experience of a bowl of salad in a restaurant. To start with, there's just a salad there. But if I look more carefully, I may see the particular leaves. If I look more carefully again, I may see the veins on the leaves and the whorls of their edges. I may start to see particular colours that I had not noticed before.

I can start to see the way the light and shade interact with those colours. If I look carefully enough, I will surely start to see types of things that I have never noticed before.

The idea, here, is that the above considerations should lead to a specific phenomenological claim about experience. Various strengths of claim are possible. The strongest claim is that experience is inexhaustibly rich – that there is always more to be found. If (perhaps for theoretical reasons) we find that implausible, we might prefer to say that experience typically, or often, contains aspects, there to be found, that go beyond anything we might now have a name for or might have thought about as such before. It is, and it should be, hard (in some sense, impossible) to convey this richness in words. Examples (recalled or, preferably, lived through) are required. The bowl of salad scarcely does the issue justice. But consider car lights reflecting on a rainy road. Consider a bright, sunlit cornfield on a windy day. Or consider whatever view you now have, even if it is of a man-made, antiseptic office. Now, stop taking the world for granted, and look again at what is there to be found. You will find (I would suggest) that experience always goes beyond what we expect, is always capable of surprising us. This is the richness in question.

The internalist conception of experience seems committed to supposing that experience, conceived of as occurring "in the head," already contains all this richness (at least in key cases, such as looking whilst not moving). Thus, it would seem, the richness of the lettuce is somehow copied into corresponding rich dynamics in the brain. Then, on this view, our conceptual thought would be able to examine the richness of the lettuce because of these rich experiential dynamics, thereby finding what was already there to be found in the internal experience but was just not yet attended to.

However, according to the externalist conception of experience that we are defending here, it is the salad itself that contains all this rich detail, and our experience – which is the right type of skilful, involved interaction with the salad – is what enables us to encounter new aspects of this richness as we attend to it.

Therefore, the internalist view can only match the phenomenal richness of experience if experience is not really as rich as it at first seems to be. If experience is rich, but not inexhaustibly rich, then a certain complex brain dynamics, caused by the world, might be an acceptable match for this. But if we really take seriously what we find when we stop taking the world for granted – that experience is inexhaustibly rich, because the world is inexhaustibly rich – then no brain dynamics can directly correspond to this encounter with worldly richness. But the ongoing, world-involving dynamics of our entire brain-body-world interaction can be this encounter with richness.

5. Conclusions

The notion of virtual action offers a good opportunity to link the ideas developed in sensorimotor and other embodied/enactive approaches with the wider goal of understanding complex mental activity, such as action-planning. Research on such questions have traditionally been the remit of top-down, cognitivist perspectives, starting from good old-fashioned AI's approach to action sequences, schemes, scripts and so on followed by less general and abstract and more situated and embodied perspectives.

Understanding virtual actions from a sensorimotor angle, therefore, is one of the general aims of the project in moving from 'simple' perception/action scenarios to a theory of cognition in general. The question that we have examined in this report is whether this meeting between sensorimotor ideas and more traditional explanatory aims necessarily involves a merging of the newer sensorimotor ideas with the more traditional explanatory strategies; in particular, with internal models and representations. Our answer is that this need not be the case.

We have defined a virtual action as an action which is not fully actualised but which can have consequences for the agent and the world. We have examined some of the existing approaches that deal with this idea and classified them based on their assumptions. Thus, many such approaches assume that since an action is virtual, most of the explanatory burden lies in internal processes, typically in the agent's brain. In such approaches, the world is only involved as a source of informational input (the same could be said of the body). We have called these 'in-the-head' approaches. They all involve various sorts of internal modelling (neural re-use of forward/inverse movement models; probabilistic hierarchical action/sensation mappings; emulator circuits; etc.).

Other approaches emphasize enabling and constitutive roles for the world and the agent's bodily coupling with it. We have called these 'world-involving' approaches. This emphasis, however, is not intended to diminish the role of processes internal to the agent, i.e., these approaches typically advocate a co-constitutional thesis by which both agent and world are required for virtual action.

While internalist approaches are nearly always – or even necessarily – representational, world-involving approaches may or may not be so. In terms of potential compatibility with SMCT, we discussed Seth's (2014) proposal of merging together sensorimotor and predictive coding ideas – an in-the-head, representational approach – and Wheeler's (2005) action-based representations – a world-involving representational approach.

Apart from these options we also discussed ecological and dynamical, non-representational, notions of strong anticipation, and argued that these already provide many of the elements for a world-involving, non-representational perspective on virtual actions. The benefits of pursuing these proposals lies in that they follow the general spirit of the sensorimotor perspective as expounded in the primary literature.

We have then proposed an explicit account of virtual actions, based on these dynamical ideas, and argued that the enactive notion of sense-making already has an implicit element of virtuality in a temporally and spatially 'thick' here-and-now: in a sense, all action is to some extent virtual action, as much as it involves an engagement of the agent with the virtual possibilities that surround the current situation in order to be normative (the risk to be avoided, the benefit to be obtained). An agent is therefore always sensitive to virtuality, i.e., to the real, contingent possibilities of the current sensorimotor situation. Understanding this sensitivity in enactive, ecological, dynamical terms, in terms of traces and tendencies in sensorimotor trajectories (as exemplified in various models) already removes the difficulty that non-representational approaches supposedly must face: how to act guided by aspects of the body and the world that have not yet been actualised. What remains as an open question is the understanding of how an agent can become increasingly sensitive to more subtle and time-extended virtual possibilities afforded by the current situation, or in ecological terms, how she may be able to 'navigate' a landscape of affordances.

In the final section of the report, we develop a complementary strand of philosophical work about virtual action. It is well-known – even notorious – that the primary SMC literature attempts to fully analyse perceptual experience in terms of the mastery of various types of sensorimotor actions. Here, we argue that this sensorimotor analysis can only be made to work if counterfactual (as well as real, or actual) actions are included. (Virtual and counterfactual actions are essentially similar ideas – we have used the term 'counterfactual' in the final section of the report, due to historical precedence in the relevant literature which we discuss there.)

There are several apparently strong objections to the SMC framework, including the argument that the approach cannot possibly account for the experience of patients with locked-in syndrome. We have argued that such objections can only be fully responded to if one is clear about the essential role of virtual action in a sensorimotor analysis of experience.

Furthermore, we have argued that this additional clarity about virtual action helps to move forward the philosophical debate on experience in other areas. In particular, we have suggested that this revised approach opens up the possibility of a genuinely scientific understanding of the idea that experience may not be just 'in-the-head'. We have argued for this by looking at existing implicit assumptions made within neurophenomenology, and criticising some of them. We propose, instead, the apparently controversial idea that perceptual experience literally involves the external objects which are perceived. This idea is not only scientifically acceptable, but fits much better with the phenomenology of experience. This approach ties in with a strong recent trend in the philosophy of perception, towards what is called direct, or naïve, realism. Simplifying somewhat, these are philosophical terms for taking seriously the idea that experience actually is as it naïvely seems to be: that what we see are objects, not representations of objects. The sensorimotor theory helps to make clear what objects *are*, such that they might be seen in this way. They are possibilities for interaction: the kind of object I can *see* is (in a way) the kind of object I can *do*. Counterfactual or virtual actions are needed here, too, in order to show how to respond to apparently strong objections to these claims about experience, based on arguments from illusion and hallucination.

References

Beaton, M. (2009). *An Analysis of Qualitative Feel as the Introspectible Subjective Aspect of a Space of Reasons*, University of Sussex D.Phil. (Ph.D.) thesis.

Beaton, M. (2013) Phenomenology and Embodied Action. *Constructivist Foundations* 8(3): 298-313

Beer, R.D. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior,* 11(4): 209-243.

Brincker, M. (2014). Navigating beyond "here & now" affordances—on sensorimotor maturation and "false belief" performance. *Front. Psychol.* 5:1433. doi: 10.3389/fpsyg.2014.01433

Buhrmann T., Di Paolo E. A., and Barandiaran X. (2013). A dynamical systems account of sensorimotor contingencies. *Front. Psychol.* 4:285. doi: 10.3389/fpsyg.2013.00285.

Buhrmann, T. and Di Paolo, E. A. (2014a). Non-representational sensorimotor knowledge. In A.P. del Pobil et al. (Eds.): *From Animals to Animats 13, Proceedings of the 13th International Conference on Simulation of Adaptive Behavior, SAB 2014*, LNAI 8575, pp. 21–31, NY, Springer Verlag.

Buhrmann, T. and Di Paolo, E. A. (2014b). Spinal circuits can accommodate interaction torques during multijoint limb movements. *Front. Comput. Neurosci.* **8**:144. doi: 10.3389/fncom.2014.00144.

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

Clark, A. (1997). *Being there: Putting brain, body, and world together again*, Cambridge, MA: MIT Press.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behav Brain Sci.*, 36(3):181-204.

Clark A. and Grush, R. (1999). Towards a cognitive robotics, *Adaptive Behavior*, 7(1), 5-16.

Clark A. and Thornton, C. (1997). Trading spaces: Computation, representation and the limits of uniformed learning, *Behavioral and Brain Sciences*, 20: 57-90.

Clark, A., & Toribio, J. (1994). Doing without representing. *Synthese*, 101, 401–31.

Colombetti, G. (2014). *The feeling body: Affective science meets the enactive mind*, Cambridge, MA, MIT Press.

Cuffari, E. Di Paolo, E., De Jaegher, H. (2014) From participatory sense-making to language: There and back again, *Phenomenology and the Cognitive Sciences*, doi 10.1007/s11097-014-9404-9 (online first).

Darwin, C. (1998/1872). *The expression of the emotions in man and animal*s. London: Harper Collins.

Degenaar, J. and Myin, E. (2014). Representation-hunger reconsidered, *Synthese*, 191: 3639-3648, doi: 10.1007/s11229-014-0484-4.

Di Paolo E. A., Barandiaran X.E., Beaton M., and Buhrmann T. (2014). Learning to perceive in the sensorimotor approach: Piaget's theory of equilibration interpreted dynamically. *Front. Hum. Neurosci.* 8:551. doi: 10.3389/fnhum.2014.00551.

Di Paolo, E. A. (2014). The worldly constituents of perceptual presence. *Front. Psychol.* 5:450. doi: 10.3389/fpsyg.2014.00450.

Di Paolo, E. A., Rohde, M. and De Jaegher, H., (2010). Horizons for the Enactive Mind: Values, Social Interaction, and Play. In J. Stewart, O. Gapenne and E. A. Di Paolo (eds), *Enaction: Towards a New Paradigm for Cognitive Science*, Cambridge, MA: MIT Press, pp. 33 – 87.

Dreyfus, H. (2007). Why Heideggerian ai failed and how fixing it would require making it more Heideggerian, *Philosophical Psychology,* 20 (2):247 – 268.

Dubois, D., (2003). Mathematical foundations of discrete and functional systems with strong and weak anticipations. *Lecture Notes in Computer Science* 2684, 110-132.

Duff, A., Fibla, M. S., & Verschure, P. F. M. J. (2011). A biologically based model for the integration of sensory-motor contingencies in rules and plans: A prefrontal cortex based extension of the Distributed Adaptive Control architecture. Brain research bulletin, 85(5), 289–304.

Friston, K. (2008). Hierarchical models in the brain. *PLoS Computational Biology*, 4:e1000211

Fish, W. (2009). *Perception, hallucination, and illusion*. New York, NY: Oxford University Press.

Froese, T. (2014). Steps toward an enactive account of synesthesia. *Cognitive Neuroscience* 5(2)

Gallagher S. & Zahavi D. (2008) *The phenomenological mind*. Routledge, Abingdon.

Grush, R. (2004) The emulation theory of representation: Motor control, imagery, and perception. *Behav Brain Sci* 27(3).

Hoffmann, M, Schmidt, N., Pfeifer, R, Engel, A., & Maye, A. (2012). Using sensorimotor contingencies for terrain discrimination and adaptive walking behavior in the quadruped robot Puppy. In T. Ziemke, C. Balkenius, & J. Hallam (Eds.), *From Animals to Animats 12* (pp. 54–64). Berlin Heidelberg: Springer.

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

Izquierdo-Torres, E., and Di Paolo, E. A., (2005). Is an embodied system ever purely reactive? *Advances in Artificial Life, Proceedings ECAL 2005, 8th European Conference*, Canterbury UK, M. S. Capcarrere et al. (eds), Springer, Berlin Heidelberg, LNAI 3630, pp 252-261.

Izquierdo, E., Harvey, I. and Beer, R. D. (2008). Associative learning on a continuum in evolved dynamical neural networks. *Adaptive Behavior,* 16:361-384.

James, W. (1968). What is an emotion? In M. B. Arnold (Ed.), *The nature of emotion: Selected readings*. London: Penguin. (Reprinted from Mind, 9, 188–205, 1884)

Kawato, M. (1999). Internal models for motor control and trajectory planning. *Curr Opin Neurobiol*, 9(6), 718–727.

Kawato, M., Kuroda, T., Imamizu, H., Nakano, E., Miyauchi, S., & Yoshioka, T. (2003). Internal forward models in the cerebellum: fMRI study on grip force and load force coupling. *Progress in Brain Research*, 142, pp. 171-188.

Kelso, J.A.S., Tuller, B., Vatikiotis-Bateson, E. & Fowler, C.A. (1984) Functionally specific articulatory cooperation following jaw perturbations during speech: Evidence for coordinative structures, *Journal of Experimental Psychology: Human Perception and Performance*, 10(6): 812–832

Kyselo, M. and Di Paolo, E. A. (2014). Locked-in syndrome: A challenge for embodied cognitive science. *Phenom. Cogn. Sci.*, doi: 10.1007/s11097-013-9344-9.

Laureys S. (ed.) (2005) *The boundaries of consciousness: Neurobiology and neuropathology*. Elsevier, Amsterdam.

Martin M. G. F. (2006) On being alienated. In: Gendler T. S. & Hawthorne J. (eds.) *Perceptual experience*. Oxford University Press, Oxford: 354-410.

Maye, A., & Engel, A. K. (2011). A discrete computational model of sensorimotor contingencies for object perception and control of behavior. *2011 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 3810–3815).

Maye, A., & Engel, A. K. (2012a). Time Scales of Sensorimotor Contingencies. In H. Zhang, A. Hussain, D. Liu, & Z. Wang (Eds.), *Advances in Brain Inspired Cognitive Systems, Lecture Notes in Computer Science* (pp. 240–249). Springer Berlin Heidelberg.

Maye, A., & Engel, A. K. (2012b). Using Sensorimotor Contingencies for Prediction and Action Planning. In T. Ziemke, C. Balkenius, & J. Hallam (Eds.), *From Animals to Animats 12, Lecture Notes in Computer Science* (pp. 106–116). Springer Berlin Heidelberg.

Maye, A., & Engel, A. K. (2013). Extending sensorimotor contingency theory: prediction, planning, and action generation. *Adaptive Behavior*, 21(6), 423–436.

McDowell J. (1996) *Mind and world*. Harvard University Press, Cambridge MA.

Merleau-Ponty M. (1962) *The phenomenology of perception*. Routledge & Kegan Paul, London. French original published in 1945.

Metzinger, T. (2014). How does the brain encode epistemic reliability? Perceptual presence, phenomenal transparency, and counterfactual richness. *Cognitive Neuroscience* 5(2)

Noë A. (2004) *Action in perception*. MIT Press, Cambridge MA.

Ostry, D. J, & Feldman, A. G. (2003). A critical evaluation of the force control hypothesis in motor control. *Exp Brain Res*, 153(3), 275–288 doi:10.1007/s00221-003-1624-0

Philipona D. & O'Regan J. K. (2006) Color naming, unique hues, and hue cancellation predicted from singularities in reflection properties. *Visual Neuroscience* 23(3-4): 331-339.

Philipona D. & O'Regan J. K. (2008) Reply to Johnson and Wright. *Visual Neuroscience* 25(2): 225-226.

Powers, W. T. (1973). *Behavior, the control of perception.* Aldine de Gruyter.

Sánchez-Fibla M., Duff A., Verschure P. F. M. J. (2011) The acquisition of intentionally indexed and object centered affordance gradients: A biomimetic controller and mobile robotics benchmark. *IEEE/RSJ International Conference Intelligent Robots and Systems (IROS),* 2011.

Sánchez-Fibla M., Duff A., Verschure P. F. M. J. (2013) A sensorimotor account of visual and tactile integration for object categorization and grasping. IEEE/RSJ International Conference on Robotics and Automation (ICRA): 107-112, DOI: 10.1109/ICRA.2013.6630563

Seth, A. K. (2014). A predictive processing theory of sensorimotor contingencies: Explaining the puzzle of perceptual presence and its absence in synesthesia. *Cognitive Neuroscience* 5(2)

Stepp, N. & Turvey, M.T. (2009) On strong anticipation, *Cognitive Systems Research*, 11(2), pp. 148–164.

Thorpe, S.J. (2002). Ultra-rapid scene categorization with a wave of spikes, *Lecture Notes in Computer Science*, 2525(1), pp. 1–15

Tuller, B., Case, P., Ding, M., & Kelso, J. A. S. (1994). The nonlinear dynamics of speech categorization. J*ournal of Experimental Psychology: Human Perception and Performance*, 20, 3–16.

Van Rooij, I., Bongers, R. M., & Haselager, W. (2002). A non-representational approach to imagined action. *Cognitive Science*, 26(3), 345–375.

Varela F. J. (1996) Neurophenomenology. *Journal of Consciousness Studies,* 3(4): 330-349.

Wallot S., and van Order, G. (2012). Ultrafast cognition, *Journal of Consciousness Studies*, 19(5-6): 141-160.

Williams, P.L., Beer, R.D., and Gasser, M. (2008). An embodied dynamical approach to relational categorization. In B.C. Love, K. McRae and V.M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 223-228).

Wheeler, M. (2005). *Reconstructing the cognitive world: The next step.* Cambridge, MA: MIT Press.

Wolpert, D M., & Miall, R. C. (1996). Forward models for physiological motor control. *Neural Networks* 9(8), 1265–-1279.

Wolpert, D. M, Doya, K. and Kawato, M. (2003) A unifying computational framework for motor control and so-cial interaction. *Philos Trans R Soc Lond B Biol Sci* 358(1431): 593–602.