

Non-representational sensorimotor knowledge

Thomas Buhrmann¹ and Ezequiel Di Paolo^{1,2}

¹ IAS Research Centre, University of the Basque Country, Spain
thomas.buehrmann@google.com

² Ikerbasque, Basque Foundation for Science, Spain

Abstract. The sensorimotor approach argues that in order to perceive one needs to first “master” the relevant sensorimotor contingencies, and then exercise the acquired practical know-how to become “attuned” to the actual and potential contingencies a particular situation entails. But the approach provides no further detail about how this mastery is achieved or what precisely it means to become attuned to a situation. We here present an agent-based model to show how sensorimotor attunement can be understood as a dynamic and non-representational process in which a particular sensorimotor coordination is enacted as a response to a given environmental context, without requiring deliberative action selection.

Keywords: sensorimotor contingencies, know-how, mastery, attunement

1 Introduction

The sensorimotor approach to perception argues that in order to perceive one must have “mastered” the relevant sensorimotor contingencies (SMCs), i.e. one must acquire a kind of practical know-how, or implicit knowledge, of the laws governing the correlation between bodily movement and associated sensory stimulation [1]. Moreover, to perceive here and now one has to exercise or deploy the mastered know-how and “tune into” the actual and potential contingencies of the current situation (ibid.). But the primary literature on the subject is mostly silent on the how this mastery is achieved, what form the practical know-how might take, or what kind of process the notion of attunement refers to. The purpose of this paper is to illustrate with a model what it might mean to exercise one’s practical know-how in order to become attuned to a situation and enact the appropriate SMCs. But since the notion of attunement is tightly linked to that of mastery we have to first discuss the relation between the two concepts. In order to develop some intuition as to how they are to be understood, we can take a look at how they are used:

“Over the course of life, a person will have encountered myriad visual attributes and visual stimuli, and each of these will have particular sets of sensorimotor contingencies associated with it. Each such set will have been recorded and will be latent, potentially available for recall: the brain thus has mastery of all these sensorimotor sets. But when a particular attribute is currently being seen, then the particular sensorimotor contingencies associated with it are no longer latent, but are actualized, or being currently made use of. [...] among all

previously memorized action recipes that allow you to make lawful changes in sensory stimulation, only some are applicable at the present moment. The sets that are applicable now are characteristic of the visual attributes of the object you are looking at, and their being currently exercised constitutes the fact of your visually perceiving that object.” (ibid., p. 945)

“[...] seeing is constituted by the brain’s present attunement to the changes that *would* occur as a consequence of an action on the part of the perceiver” (ibid., p. 968, italics added)

It is clear that “mastery” is supposed to refer to the “accommodation” of certain regularities in the environment, and attunement to the exploitation of these regularities. But “accommodation” in this context does not necessarily mean that the contingent aspects of the environmental regularities are stored internally by the agent, but simply that the agent has undergone some changes such that whenever the regularities present themselves in a new situation, the agent is able to re-enact sensorimotor engagements that have previously been adequate in similar sensorimotor situations. It is also implied that the act of exercising one’s SM knowledge is not a deliberative process of consciously weighing different possible SM coordinations to engage in. It rather seems to be an automatic process in which the right coordination is solicited as a response to a particular situation, in other words a kind of “resonance” between environment and agent. Moreover, from the second quote it follows that the exercise of my know-how can be counterfactual, i.e. my perception of possibilities for interaction with an object depends not only on my current engagement with it, but also on my practical knowledge of properties of the object that are not directly available.

Before describing a dynamic model that we think captures the essence of the process of attunement, we draw on Merleau-Ponty’s account of skill acquisition [2–4] to further elucidate some of these notions. Three aspects characterise the learning of sensorimotor skills according to Merleau-Ponty. Firstly, in the acquisition of everyday skills, the accumulation of experience serves to discriminate situations that solicit a particular response with increasing specificity. Secondly, experience also allows a person to incrementally refine her dispositions to respond to these solicitations. And thirdly, behavioural responses to a situation take the form of movement towards the completion of a Gestalt, or equilibrium, to which the body tends without the need to mentally represent it. Though still a rather abstract account of skill acquisition, translating these three elements into dynamical systems terminology allows us to arrive at a description that will be useful in the analysis and interpretation of models addressing the issues of mastery and attunement.

In correspondence to the three aspects, firstly, if we consider an agent as a dynamical system coupled to its environment, then different environmental conditions, reflected in different sensory inputs, can result in the divergence of initially identical agent states. At a future point in time, therefore, the agent can react to the same sensory stimulus in different ways, as the accumulated history of its coupling with the environment has left the agent in different parts of its state space. In other words, the accumulation of experience allows the agent to discriminate between different contexts when exposed to identical sensory perturbations. Secondly, the behaviour of the agent as a dynamical

system depends on its limit sets. Since through continued environment interaction the agent is able to reach different areas of state space, from these different initial conditions the agent may then follow different behavioural tendencies as determined by its attracting and repelling sets. Experience therefore can also serve to tune those limit sets such that the agent’s movement through state space corresponds to the desired response that a given situation solicits. And thirdly, the agent’s movements are fully determined by the relaxation of its dynamics towards the limit set in whose basin it finds itself at any given time. And the agent can not in any meaningful way be said to represent what the final state is that it is tending towards.

In short, the learning of a SM skill, in this view, corresponds to the tuning of the agent’s dynamical landscape such that different environmental contexts leave the agent in different parts of state space, and such that the appropriate response corresponds to a particular trajectory of the dynamics when relaxing towards equilibrium. We would like to suggest here then, that in the process of mastery an agent’s dynamical landscape is shaped over time to incrementally refine the discrimination of and response to different environmental situations; and that attunement is a process of interaction with the environment, such that a particular situation solicits the appropriate sensorimotor coordination.

We next describe a model to illustrate this latter interpretation of attunement in more detail.

2 Materials and Methods

The model presented here consists of an agent artificially evolved to identify only through touch the properties of a planar surface presented at different relative orientations and positions. After a period of unconstrained interaction with the surface, the agent has to demonstrate that it has retained something about its properties by producing movements following the orientation of the surface without the surface being present any longer (i.e. without the corresponding sensory stimulation). In other words, the agent has to re-enact the now invisible surface, or act as if it was still present. Note that we consider the evolved agent as already having mastered the required skill. The analysis of the agent’s behaviour focuses on how the acquired know-how is exercised.

The model is summarised in panel A of Figure 1. The agent’s body is a two-joint arm controlled by a continuous-time recurrent neural network [5] and equipped with a touch sensor (pink, dotted line). The environment consists of a planar surface (green line) whose position and orientation relative to the arm can vary. The agent’s touch sensor can register the distance to the surface when in close proximity, but the arm can freely pass through it. The agent’s neural network has a fully connected hidden layer receiving three inputs: the arm’s two joint angles $\theta_{1,2}$ and the distance D between end-effector and surface as measured by the sensor. Two output neurons are fully connected to the hidden layer and control the desired joint angles $\theta_{1,2}^d$. These are transformed by PD controllers into joint torques that are applied to the arm (blue arrows) to produce the required movement. The arm dynamics are given by a common model derived from Newton-Euler equations and d’Alembert’s Principle (for details see [6]). The parameters of the PD controllers were tuned by hand to achieve a somewhat underdamped response.




Fig. 1. Experimental setup. A: two-joint planar robot arm controlled by PDs whose set points (desired joint angles $\theta_{1,2}^d$) are determined by the outputs of a recurrent neural network. B: Dimensions of the arm and range of surface positions and orientations. Surfaces are presented at two positions and orientations covering a range of 30 degrees. C: time course of each trial showing the progression of evaluation phases and the state of the sensor.

All nodes in the agent's neural network are modelled as leaky integrators:

$$\tau_i \dot{y}_i = -y_i + \sum_{j=1}^n w_{ji} \sigma(y_j + \theta_j)$$

where y_i is the activation of neuron i , $\tau_i \in [0.01, 4]$ its time constant, $w_{ji} \in [-10, 10]$ the strength of the connection from neuron j to i , $\theta_j \in [-10, 10]$ a bias term, and $\sigma(x) = 1/(1 + e^{-x})$ a sigmoidal activation function. Both arm and neural dynamics are Euler integrated with a step size of 0.05.

The surface can vary in position and angle, and a unique combination is tested in each experimental trial. The range of surfaces evaluated is shown in panel B of Figure 1. The arm (shown in black) at maximum extension has a length of 0.65 units (indicated by the light grey half-disk). The distance sensor is attached at the end of the arm and can sense objects up to 0.05 units ahead (dark grey half-disk). Its response signal is inversely proportional to the sensed distance and scaled to the range $[0, 1]$. Surfaces

are presented at two different positions (0.44 and 0.56 units from the arm’s origin) and are 1.4 units in length. The angular range covered by the surfaces is 30° ($0 \pm 15^\circ$). Two more surfaces per position are used in the experiment (at $\pm 7.5^\circ$) but not shown here for clarity.

The task is to move the agent’s end effector along the particular surface presented in each trial, even after its touch sensor is disabled. A version of the microbial genetic algorithm [7] is used to search for neural network parameters that allow the agent to solve this task. Each candidate solution is evaluated on 10 trials, in each of which the agent is presented with a different surface (2 positions x 5 orientations). The time course of fitness evaluation in each trial is shown in panel C of Figure 1. During the first 2 seconds the agent’s movements are unconstrained, i.e. its behaviour does not contribute to its measured fitness. During the “probe” phase (2.6 s), the agent is rewarded for proximity to the surface but is otherwise free to move in an arbitrary manner. Its fitness in this period is equal to the end effector’s average proximity to the surface (measured using the shortest distance). In the “follow” phase (last 4 seconds), the agent’s fitness is determined by the combination of average proximity and the end effector’s average velocity parallel to the surface (measured as the absolute length of the projection of the velocity vector onto the surface). After blackout, the agent can no longer sense the surface, yet has to keep moving along it (e.g. unidirectional or by oscillating back and forth in the corresponding plane). The total fitness of the agent in a single trial is the average of the fitness achieved in the probe and follow periods, and the overall fitness across all trials is equal to the minimum of individual trial fitnesses.

Given this experimental setup, there are two types of solutions to the task. After blackout, i.e. when the agent is “touch-blind” and has to re-enact the previously encountered surface orientation, the only sensory inputs available are the agent’s current joint angles. Hence, if the agent were to rely on sensory inputs only to discriminate the different surfaces and produce a different SM coordination in response, this would imply that the joint configurations at the time of blackout would have to be unique for all surfaces encountered. If, in contrast, the joint configurations at blackout are not unique, then a successful agent must have used the initial exploration phase to reach different parts of its state space, such that different behaviours can be produced in response to identical sensory states. Such a process of state differentiation in covariation with relevant environmental variables we would then be happy to label “attunement”.

3 Results

Successful agents evolved reliably with as few as 3 hidden neurons, but better performance could be achieved with a larger number. In the following we present results for an agent with 8 hidden neurons, which achieved a fitness of 82%.

3.1 Evolved behaviour

In panels A-D of Figure 2 we show typical examples of the agent’s end effector trajectories (grey, darker shades indicating greater touch response) overlaid on top of each corresponding environmental surface (red). The data is shown in joint space (θ_2 vs. θ_1),

i.e. the red lines correspond to those joint configurations the arm would have to adopt to reach points on the surface. The left and right column show data for the two extremal surface orientations. The first and second row correspond to surfaces at the closer and farther distance respectively. Note that the simple planar environmental surfaces are in fact complex curves in motor space; and that surface orientation and position in Cartesian space seem to become surface position and scale in motor space respectively.




Fig. 2. Joint space trajectories of successfully evolved agent. A-D: Shown in red are those joint configurations that the agent would have to adopt to reach points on the surface, i.e. the planar surface translated into joint space. Overlaid in grey are the performed trajectories, with darker shades indicating greater touch sensor activity. Green markers indicate the initial position. A and B correspond to surfaces at different orientations (± 15 degrees). C and D correspond to the same orientations as A and B respectively, but at a greater surface distance ($+ 0.12$). E: Trajectories for three surface positions (in red, green and blue) and five orientations, darker shades again indicating greater touch response.

Inspecting the trajectories we observe that the agent initially performs a stereotypical transient that eventually makes contact with each surface at a point that depends on the configuration of the surface (the trajectories turn from grey to black here). All points of first contact occur at the "lower" end of the area of highest curvature. After contact has been established, the agent uses its sensor to move the end effector along the surface in a single direction. Finally, when the sensor is disabled (trajectories turn grey again), the agent continues to move along the opposing side of the surface. While towards the end of the trial the agent begins to deviate from the surface, for a significant time after blackout the agent manages to follow it well.

At first glance the observed behaviour might not seem remarkable. However, the agent has control only over the position of the arm in joint space. Hence to follow the curvature of the surface, the neural network has to produce a complex time-series of joint angles such that the corresponding end-effector positions lie on that same curve.

In other words, after blackout the agent cannot rely on some form of inertia to keep it moving along the required curve. Nor can the agent simply hold certain variables fixed in order to keep moving in the same direction, as would be the case for a wheeled robot following a straight line for example.

The most important aspect of the observed behaviour is the fact that for each surface the agent produces a different trajectory after blackout, even though the touch sensor returns the same signal in all cases. This difference in SM behaviour in response to different surfaces becomes more salient when we draw all trajectories in a single figure, as shown in panel E for three different surface positions and five orientations each. We observe that some end effector trajectories cross each other in motor space but subsequently follow different paths. Since during this phase in the trial the same joint angles also serve as the only inputs to the agent’s neural network (motor and sensory space are the same), it is already clear that the behaviour cannot be determined by instantaneous sensory input alone (since the same sensory input here leads to different behavioural responses). Instead, the differentiation of identical sensory configurations must be based on the history of the agent’s engagement with the surface, as accumulated in the agent’s state.

3.2 Generalization

Before identifying what sort of mechanism underlies the agent’s unique behavioural response to identical sensory stimulations, we will demonstrate how the agent’s performance generalises to a larger range of surface configurations than encountered during evolution, and identify the kind of features the agent might use to achieve this.

In panel A of Figure 3 we plot a heatmap of the agent’s fitness as a function of surface orientation (horizontal axis) and position (vertical axis) at 30 different values each. The surface configurations used during evolution are located at the intersections of the red lines. It is clear that the agent establishes successful interactions with the surface over a wide range of surface configurations.

This task would be relatively simple if there existed a simple, instantaneous sensory feature that uniquely identifies each of the 900 tested surfaces. For example, since the agent initially performs a stereotypical transient, one could imagine that the time of first contact with the surface might be unique, or equivalently the joint configuration at this point in time; or, perhaps, the joint angles at the point when the sensor is disabled.

To see whether this is the case, in panel B of Figure 3 we plot for each of the 900 surface configurations a point whose coordinates are the time of first contact (t_c), as well as the state of the two joint angles when the sensor is disabled (denoted by $\theta'_{1,2}$). One can observe that the resulting points lie on a curved surface that folds back on itself in the t_c dimension. From the curvature of this manifold two implications can be derived. Firstly, there is no unique time of first contact across all surfaces (e.g. the top part of the manifold is curved such that points in closer and farther regions relative to the plane of the screen can have identical t_c coordinates). This can in fact be seen already in panel E of Figure 2, where trajectories for different surfaces split from the initial transient at the same time. It follows that the state of the proprioceptive joint sensors at this time cannot be unique. Secondly, because of the fold, the joint angle configuration $\theta'_{1,2}$ when the sensor is disabled does not uniquely identify each surface either.




Fig. 3. Performance for 900 different surface configurations (30 orientations x 30 positions). A: heatmap of fitness achieved for each surface (for fitness scale see colour bar shown above, theoretical maximum = 1.0). The intersections of red lines indicate the 10 surfaces encountered during evolution. The dark area on the right corresponds to surfaces that the agent failed to re-enact. B: Scatter plot showing for each of the 900 surfaces a 3d point whose coordinates are the time of first contact between end effector and surface (t_c), and the joint configuration of the arm when the touch sensor is disabled ($\theta'_{1,2}$).

What is unique, however, is the combination of these two features, as there are no two points that coincide in both t_c and $\theta'_{1,2}$. It is clear that there must be some combination of features that can be used to distinguish all surfaces, as otherwise the agent would not be able to solve the task. The question is how the agent manages to retain information about the initial contact such as to respond appropriately later on when the sensor is disabled. In other words, how does the agent integrate sensory signals over time such that it can respond uniquely to ambiguous sensory information at the time of sensor blackout?

3.3 Neural mechanism underlying surface disambiguation

To better understand how the agent differentiates between surfaces even when instantaneous sensory feedback is not unique we look at the agent's dynamics in higher dimensions. Figure 4.A shows a projection of the agent's state that consists of the two joint angle sensors ($\theta_{1,2}$) and the output of a hidden neuron (h_1). Lines of the same colour correspond to trials where surfaces are positioned at the same distance but oriented differently. Colour shade varies with orientation, and different colours correspond to different surface positions. Black dots mark the time when the sensor is disabled.

Starting from a common initial position, the combination of all initial transients forms an arc that traces the stereotypical movement pattern exhibited by the agent if no surface was present. For different surfaces, then, the trajectories separate from this arc at different points, namely the point at which contact is made with the surface. One can observe that trajectories belonging to surfaces of different orientations but identical




Fig. 4. 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.

distances (i.e. those of the same colour) are already ordered at this time. In other words, the time of first contact correlates with surface orientations.

However, there are also trajectories that start diverging at the same point, even though the corresponding surfaces are located in different positions (differently coloured lines); i.e. trajectories belonging to different surface positions are still partly “entangled”. During the period of sensor-based surface interaction, however, these trajectories eventually become disentangled, forming clearly separated bundles. For each surface position corresponding trajectories now form curved manifolds that lie parallel to each other, and in each of which trajectories are ordered by surface orientation. Thus when viewed from certain angles or in certain projections, trajectories belonging to different manifolds might seem to cross (like in Figure 2), while in fact being well separated in higher dimensions. Finally, just before the sensor is disabled, each manifold is being twisted. This preserves the established ordering, but ensures that subsequent parts of the trajectories are shaped such that their projections into Cartesian space are approximately straight (like the corresponding surface).

What keeps the trajectories separate after the sensor is disabled? Do they relax into different steady-states? To answer this question we observe the dynamics beyond the duration of the trial with the sensor remaining off. This makes the agent an autonomous system, which should eventually reach a steady state. The result can be seen in panel B of Figure 4, which shows the same projection as panel A, but from a different angle and for clarity only two surface positions. One can see that within the plotted duration of time, the agent approximately reaches steady-state, and in particular the same stable

attractor for all trajectories. In principle the meticulous ordering of trajectories within and across manifolds could also have been the result of trajectories tending towards different attractors. Here, however, the different sensorimotor coordinations are formed by different transients within the same attractor basin, but in such a way that their separation is maintained until the end of the trial. Similar forms of state-determined sensitivity to the environment have been studied using information-theoretic techniques [8], which could be of use in this case as well.

In summary, underlying the act of distinguishing and responding differently to various surface positions and orientations is the integration over time of several aspects of the interaction. We have seen that surfaces that are positioned identically but oriented differently can be identified in the neural dynamics by the time of first contact alone. But at this point there still exists ambiguity with surfaces located at other distances. The ongoing interaction with the surface helps to disambiguate these cases by splitting trajectories into different manifolds. All trajectories now being perfectly separated and ordered in state space allows them to relax towards a common attractor along unique transients, with the shape of transients corresponding to the desired surface orientation.

4 Discussion

Our purpose has been to illustrate via a minimal model what it means to become attuned to a situation and enact the appropriate SMCs. We have shown that attunement can be interpreted as a continuous process of agent-environment interaction such that different situations are discriminated via the separation of the agent’s dynamics. As a result, unique behavioural responses can be enacted in the form of particular transients during the relaxation towards equilibrium. Mastery of SMCs, then, corresponds to the process that shapes the dynamical landscapes of an agent such that attunement is possible.

As noted in the introduction, the dynamical process can be interpreted as a minimal example of Merleau-Ponty’s motor intentionality, and thereby illustrate how agents exhibit embodied “purposeful” behaviours, without representing environmental features or explicit goals. Just like, for example, a tennis player who performs his serve or return without having to contemplate each required step, nor the details of the ideal posture to adopt when hitting the ball, the performance of a SM skill is more like a habitual relaxation towards an optimal movement “gestalt” that is only implicit in agent-environment dynamics, and which is solicited by a certain situation in the world.

Our model does not introduce new dynamical phenomena that we suppose to underlie the process of attunement. State retention and differentiation are well-known aspects of dynamical systems. We also do not advocate the particular methodology adopted here over other alternatives for investigating SMCs, such as information-theoretic analysis [8]. Nor did we aim to add to the already extensive catalogue of minimally cognitive behaviours that even simple dynamical agents can exhibit. The purpose, rather, was to clarify a core concept of the sensorimotor approach to perception, namely that of attunement. We believe that the methodology has been adequate, and the results sufficient, to show that the selection and exercise of SMCs can be understood as a dynamical process that does not require a dedicated mechanism or organisational level at which SMCs are “represented” (in a non-trivial sense of the term), nor the invocation of

levels of explanation other than that of sensorimotor relations. Though the different SMCs enacted by our agent were solutions to variants of the same task, this does not limit the applicability of our interpretation (for enactment of radically different SMCs see e.g. [9]). Equally, while for the investigated task successful agents' dynamics transiently separated into different manifolds depending on the environment, other manners of separating and organizing dynamics (e.g. by attractor basin) are also compatible with our account of attunement.

We have not here provided a model of how SM skills are mastered, i.e., how the dynamical landscape of the agent is altered through experience (the evolutionary search employed here to identify an appropriate agent is not meant to model the process of mastery). Yet we believe that simple models of the kind presented can contribute to filling in the gaps in sensorimotor theory. Not only by making more explicit what we mean when we talk about notions such as “mastery” or “attunement”, but also by deriving implications that only become clear when these notions are operationalised [10]. One such implication is that sensorimotor theory does not have to evoke explicit representational vehicles, nor deliberative processes of action selection to account for the acquisition and exercise of SMCs. This lends further evidence to a radical reading of sensorimotor theory [11], which rejects the role of contentful representations.

Acknowledgements: Thanks to Randall Beer and Eduardo Izquierdo for their comments on an early draft of this work. This work is funded by the eSMCs: Extending Sensorimotor Contingencies to Cognition project, FP7-ICT-2009-6 no: 270212.

References

1. O'Regan, J., Noë, A.: A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences* **24** (2001) 939–1031
2. Merleau-Ponty, M.: *Phenomenology of Perception*. 2 edn. Routledge (2002)
3. Dreyfus, H.L.: Intelligence without representation. Merleau-Ponty's critique of mental representation. *Phenomenology and the Cognitive Sciences* **1**(4) (2002) 367–383
4. Kelly, S.D.: Merleau-Ponty on the body. *Ratio* **15**(4) (2002) 376–391
5. Beer, R.D.: On the dynamics of small continuous-time recurrent neural networks. *Adapt. Behav.* **3**(4) (1995) 469–509
6. Hollerbach, M.J., Flash, T.: Dynamic interactions between limb segments during planar arm movement. *Biological cybernetics* **44**(1) (1982) 67–77
7. Harvey, I.: The microbial genetic algorithm. In: *Proceedings of the 10th European conference on Advances in artificial life. ECAL'09*, Berlin, Heidelberg, Springer-Verlag (2011) 126–133
8. Beer, R., Williams, P.: Information processing and dynamics in minimally cognitive agents. To appear in *Cognitive Science*. (2014)
9. Izquierdo, E., Buhrmann, T.: Analysis of a dynamical recurrent neural network evolved for two qualitatively different tasks: walking and chemotaxis. In Bullock, S., et al., eds.: *Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*, MIT Press, Cambridge, MA (2008) 257–264
10. Buhrmann, T., Di Paolo, E.A., Barandiaran, X.E.: A dynamical systems account of sensorimotor contingencies. *Frontiers in Cognition* **4**(285) (2013)
11. Hutto, D.D., Myin, E.: *Radicalizing enactivism: basic minds without content*. MIT Press, Cambridge, Mass. (2013)