

Geometric Design Principles for Brains of Embodied Agents

Nihat Ay^{1,2,3}

Received: 9 February 2015 / Accepted: 11 June 2015 / Published online: 22 July 2015
© The Author(s) 2015. This article is published with open access at Springerlink.com

Abstract I propose a formal model of the sensorimotor loop and discuss corresponding extrinsic embodiment constraints and the intrinsic degrees of freedom. These degrees constitute the basis for adaptation in terms of learning and should therefore be coupled with the embodiment constraints. Notions of sufficiency and embodied universal approximation allow us to formulate principles for such a coupling. This provides a geometric approach to the design of control architectures for embodied agents.

Keywords Sensorimotor loop · Embodiment · Information geometry · Cheap design · Universal approximation

1 Introduction

Within the last few decades, it has become clear that an understanding of intelligence and cognition has to take into account the fact that behaving agents are embodied and situated [9, 19]. At first sight this might appear like an obvious and unimportant observation. After all, it is obvious that agents *do* have bodies and that they *are* situated in some environment. So what? But the more we think about

it the more we should understand that this observation has actually far-reaching implications. Any behaviour of an agent is mediated through its sensorimotor apparatus and its interactions with the environment. Therefore, the control of intelligent behaviour is inseparably intertwined with the agent's sensorimotor constraints. Such a coupling should allow us to reveal design principles for brains and their implemented control mechanisms. In this regard, it has been pointed out that quite complex behavioural patterns do not necessarily require complex control, leading to the notions of morphological computation and the principle of cheap design [19, 24, 25].

In this article, I want to formally address the tight connection between the control architecture and the embodiment of an agent in terms of geometry, in particular information geometry [1, 3]. In order to do so, I apply a theory of the sensorimotor loop in terms of a causal model, as developed in [7, 13, 16], and propose an approach for designing controller architectures that utilise the sensorimotor constraints. My aim is not to provide the most general results, but to exemplify the strength of the theory by integrating initial results from an existing body of work [2, 4, 8, 12, 14–16, 25], and discussing them, for the first time, from a unifying perspective. The main intuition for connecting control architectures to sensorimotor constraints is based on two assumptions. On one hand, the architecture should be sufficient for the expression of desired behaviours, given the embodiment constraints. This clearly requires some amount of richness or complexity of the architecture. On the other hand, it should be concise in some natural sense. I discuss various versions of these two assumptions and derive related results. In particular, I address the notion of embodied universal approximation and contrast it with the standard notion of universal approximation. This is closely related to the recent work

✉ Nihat Ay
nay@mis.mpg.de

¹ Max Planck Institute for Mathematics in the Sciences,
Inselstraße 22, 04103 Leipzig, Germany

² Department of Mathematics and Computer Science,
Leipzig University, 04009 Leipzig, Germany

³ Santa Fe Institute, 1399 Hyde Park Road, Santa Fe,
NM 87501, USA

[16] which was initiated by the general line of research sketched in the present article. Being based on stronger assumptions, on the other hand, the results of [16] are more refined than the corresponding ones presented in this broad treatment.

Section 2 provides the conceptual and formal definition of the sensorimotor loop of an embodied agent. To the agent, the world appears as a black box which contains the agent's body and its environment. Sensorimotor mechanisms and their generation of behaviour, which takes place in the world, are formalised. Section 3 discusses and formalises extrinsic and intrinsic sensorimotor constraints. Intrinsic constraints are given in terms of a controller model, which allows the agent to adapt to the extrinsic constraints. The notions of sufficiency and embodied universal approximation are introduced as basis for a geometric design of controller models. Finally, Sect. 4 exemplifies the developed geometric methods in the context of policy models and derives four examples of embodied universal approximators.

2 A Formal Model of Embodied Agents

2.1 The Basic Components of the Sensorimotor Loop

What is an embodied agent? In order to develop a theory of embodied agents that allows us to cast the core principles of the field of embodied intelligence into rigorous theoretical and quantitative statements, we need an appropriate formal model. Such a model should be general enough to be applicable to all kinds of embodied agents, including natural as well as artificial ones, and specific enough to capture the essential aspects of embodiment. How should such a model look like? First of all, obviously, an embodied agent has a body. This body is situated in an environment with which the agent can interact, thereby generating some behavior. In order to be useful, this behavior has to be guided or controlled by the agent's brain or controller. We draw the boundary between the brain on one side and the body, together with the environment, on the other side. We assume that the brain receives sensor signals from and sends effector or actuator signals to the outside world (body and environment), which is the only mechanism of signal transmission in both directions. Stated differently, we assume that the brain is causally independent of the world, given the sensor signals, and the world is causally independent of the brain, given the actuator signals. We refer to this perspective as the *black box perspective*. In particular, the boundary between the body and the environment is not directly “visible” for the brain and

has to be identified or actively constructed as result of the interaction with the world.

2.2 The Mechanisms of the Sensorimotor Loop

Let us now develop a formal description of this sensorimotor loop. We denote the set of world states by \mathcal{W} . This set can be, for instance, the position of a robot in a static 3D environment. Information from the world is transmitted to the brain through sensors. Denoting the set of sensor states by \mathcal{S} , we model the mechanism of a sensor in terms of an information transmission channel from \mathcal{W} to \mathcal{S} as it is defined in information theory. More precisely, given a world state $w \in \mathcal{W}$, the response of the sensor can be characterized by a probability distribution of possible sensor states $s \in \mathcal{S}$ as result of w . For instance, if the sensor is noisy, then its response will not be uniquely determined. If the sensor is noiseless, that is deterministic, then there will be only one sensor state as possible response to the world state w . In any case, the response of the sensor given the world state w can be described as a distribution $\beta(w)$ on the sensor states s . Therefore, the mechanism of the sensor can be formalised as a Markov kernel

$$\beta: \mathcal{W} \longrightarrow \Delta_{\mathcal{S}},$$

where $\Delta_{\mathcal{S}}$ denotes the set of distributions on \mathcal{S} . Each distribution $\beta(w)$ is uniquely determined by its values $\beta(w; ds)$ for infinitesimal sensor state sets ds . From this, we can calculate the probability of any (measurable) set $\mathcal{S}' \subseteq \mathcal{S}$ by integration:

$$\beta(w; \mathcal{S}') = \int_{\mathcal{S}'} \beta(w; ds).$$

Given a reference measure $\mu_{\mathcal{S}}$ on the set of sensor states so that all distributions $\beta(w)$ are continuous with respect to $\mu_{\mathcal{S}}$, we write $\beta(w; s)$ for the corresponding density of $\beta(w)$, that is $\beta(w; \mathcal{S}') = \int_{\mathcal{S}'} \beta(w; s) \mu_{\mathcal{S}}(ds)$. In particular, this always holds in the case of finitely many sensor states where we take as reference measure the counting measure. Finally, we denote the set of all sensor mechanisms by $\Delta_{\mathcal{S}}^{\mathcal{W}}$.

After having described in detail the mathematical model of a sensor, it is now straightforward to consider corresponding formalisations of the other components of the sensorimotor loop in terms of Markov kernels, where we apply, without explicitly stating, the same notation conventions as for the sensor. Based on our conceptual perspective, which is illustrated in Fig. 1b, we compose the sensorimotor mechanisms according to a causal diagram, which represents the interaction process of the agent with the world (see Fig. 2).

We now continue with the notion of a policy. The agent can generate an effect in the world in terms of its actuators.

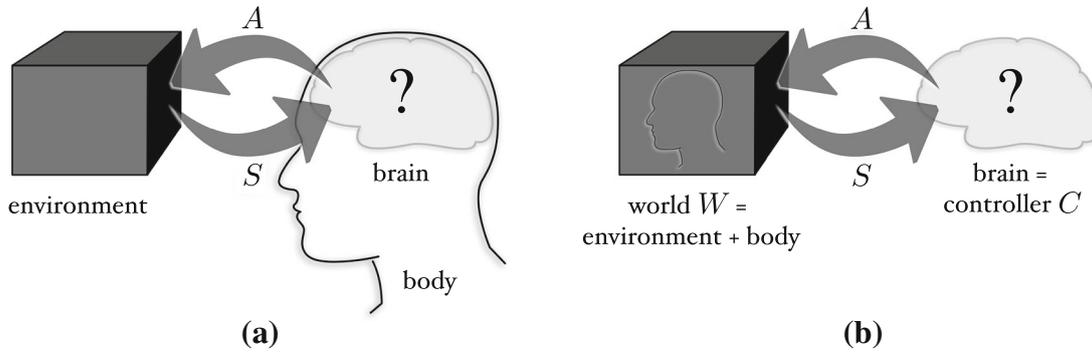


Fig. 1 Two ways to define the agent separation: **a** agent (= brain + body) interacts with its environment considered as a *black box*; **b** agent (= brain) interacts with the world (= environment + body) considered as a *black box*

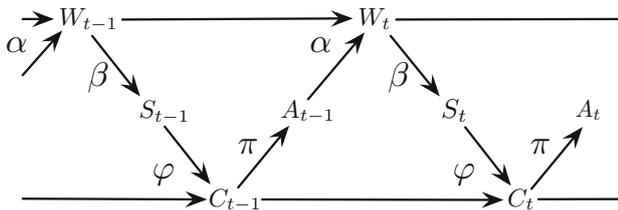


Fig. 2 The causal diagram of the sensorimotor loop as conceptually illustrated in Fig. 1b

Since we consider the body as being part of the world, this can lead, for instance, to some body movement of the agent. In order to control or guide this movement, it is beneficial for the agent to choose its actuator state based on the internal controller state which contains information about the world received through its sensors. Denoting the controller state set by \mathcal{C} and the actuator state set by \mathcal{A} , we can again consider a channel from \mathcal{C} to \mathcal{A} as formal model of a policy, which we denote by π . Being more precise, a policy receives a controller state c and generates, based on c , a distribution $\pi(c)$ of actuator states. Again, we have a Markov kernel

$$\pi : \mathcal{C} \rightarrow \Delta_{\mathcal{A}}.$$

Note that this definition of a policy allows us to also consider a random choice of actions, so-called non-deterministic policies.

Finally, we formalise the dynamics α of the world and the dynamics φ of the controller, in terms of Markov kernels

$$\alpha : \mathcal{W} \times \mathcal{A} \rightarrow \Delta_{\mathcal{W}} \quad \text{and} \quad \varphi : \mathcal{C} \times \mathcal{S} \rightarrow \Delta_{\mathcal{C}}.$$

2.3 From Mechanisms to Embodied Behaviours

In order to describe the interaction process between the agent and the world, we have to sequentially apply the

individual mechanisms in the right order. We first consider only one step as shown in Fig. 3, and then iterate this step. Starting with an initial world state w and an initial controller state c , we have the following Markov kernel that describes the transition to the new world and controller states:

$$\begin{aligned} \mathbb{P}(w, c; dw', dc') &= \int_{\mathcal{A}} \int_{\mathcal{S}} \varphi(c, s'; dc') \beta(w'; ds') \alpha(w, a'; dw') \pi(c; da') \\ &= \int_{\mathcal{S}} \varphi(c, s'; dc') \beta(w'; ds') \int_{\mathcal{A}} \alpha(w, a'; dw') \pi(c; da') \\ &= \mathbb{P}(c, w'; dc') \mathbb{P}(w, c; dw'). \end{aligned} \tag{1}$$

Iterating this one-step transition T times defines a joint process $(w^1, c^1), \dots, (w^T, c^T)$ of the world and the controller, conditioned on the initial joint state (w^0, c^0) . Behaviour is a process that takes place in the world, for instance as a particular movement of the agent’s body which we considered to be part of the world (see Fig. 1). This implies that only the world process w^1, \dots, w^T will be of relevance for the study of behaviour. Marginalising out the controller process c^1, \dots, c^T leads to

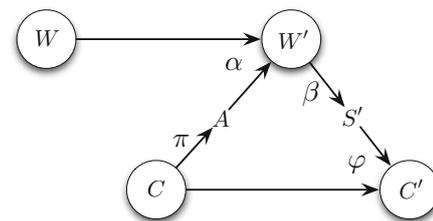


Fig. 3 The one step mechanism of going from w, c to w', c' as defined by (1)

$$\begin{aligned} \mathbb{P}(w^0, c^0; dw^1, \dots, dw^T) &= \int_{\mathcal{C}} \dots \int_{\mathcal{C}} \mathbb{P}(w^0, c^0; dw^1, dc^1, \dots, dw^T, dc^T) \\ &= \int_{\mathcal{C}} \dots \int_{\mathcal{C}} \prod_{t=1}^T \mathbb{P}(w^{t-1}, c^{t-1}; dw^t, dc^t), \end{aligned} \tag{2}$$

where we integrate T times. This Markov kernel encodes all information that is required for the evaluation of behaviour.

3 Intrinsic and Extrinsic Sensorimotor Constraints

In the previous section, I have introduced the sensorimotor loop in terms of the mechanisms π , α , β , and φ . Here, the mechanisms α and β are extrinsic to the agent and define its embodiment constraints. I refer to these constraints as being of the first kind in order to distinguish them from other constraints introduced below. Clearly, α and β are not fixed and can evolve in time. For instance, if an agent moves from one environment to another the mechanism α will change. Both mechanisms, α and β , will change, if the agent’s body is subject to a developmental process or to a partially disabling injury. However, we can expect that, due to invariant structures and regularities of the environment and the body, the mechanisms will remain in a restricted domain. This fact implies a second kind of embodiment constraint. In addition to the constraints of the first kind encoded by a single pair $\sigma := (\alpha, \beta)$ of extrinsic mechanisms, the evolution of σ is restricted to trajectories within a subset Σ of possible mechanisms, corresponding to the variability of the body and the environment. We consider the direct cause of this evolution within Σ to be completely exogenous with respect to the agent and, therefore, it can not be directly controlled by the agent. In contrast, we assume that the intrinsic mechanisms π and φ , summarised as $\gamma = (\pi, \varphi)$, can be endogenously modified by the agent in terms of a learning process. The question is whether this intrinsic process is also constrained. Obviously, this will depend on the degrees of freedom of the controller or brain architecture. If it is rich enough to implement any π and φ then we have to consider the set

$$\Sigma \times \Gamma, \quad \text{where } \Gamma := \Delta_{\mathcal{A}}^{\mathcal{C}} \times \Delta_{\mathcal{C}}^{\mathcal{C} \times \mathcal{S}}, \tag{3}$$

as domain for the process of all sensorimotor mechanisms. But should we assume this richness of brain architectures? I argue that it is not only not required but also beneficial for the agent to have a restricted control architecture, which is in line with the idea of cheap design within the field of embodied intelligence. However, the theoretical study of unrestricted control architectures, referred to as *universal approximators*, is helpful for exploring the ones that are

optimal for the given constraints Σ of the agent. Here, optimality involves two requirements. On one hand, the control architecture should be sufficient in the sense that it enables the agent to adapt to the embodiment constraints Σ . On the other hand, it should be concise in order to efficiently implement this adaptation. The field of embodied intelligence offers many case studies as evidence for such kinds of cheap control. This field highlights, in particular, the fact that quite complex and useful behaviours do not require much control [19].

Now, I am going to make these ideas on optimal architectures more precise. A brain architecture can be associated with a set $\mathcal{M} \subseteq \Gamma$, the set of those $\gamma = (\pi, \varphi)$ that can be generated by the architecture. We refer to such a set as *controller model* or simply *model*. Most generally, a model can be any subset \mathcal{M} of Γ . Typically, a model \mathcal{M} is defined as the image of a map η , referred to as a *parametrisation* of \mathcal{M} , from a parameter set $\Theta \subseteq \mathbb{R}^d$ to Γ . Figure 4 illustrates the important parametrisation in terms of synaptic couplings w_{ij} of neurons i and j . We call a model \mathcal{M} together with a parametrisation η a *parametrised model*. Clearly, any map $\eta : \Theta \rightarrow \Gamma$ is a parametrisation of the model given by its image.

In order to be useful in applications, a parametrised model is typically assumed to have further properties which are context dependent, for instance smoothness properties up to some order.

In addition to the embodiment constraints Σ , a model \mathcal{M} encodes the intrinsic constraints of the brain and represents, at the same time, the domain in which learning processes can take place. The set (3) of sensorimotor loops then reduces to

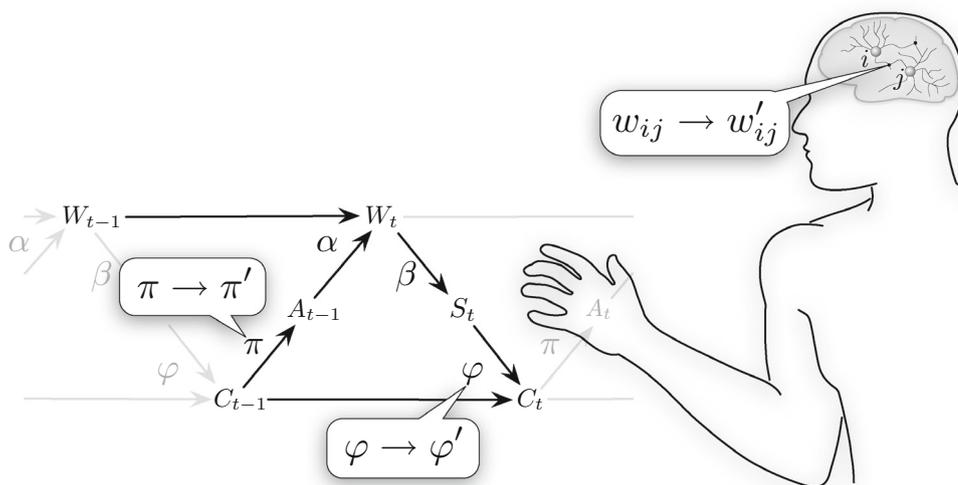
$$\text{SML} := \Sigma \times \mathcal{M}. \tag{4}$$

In what follows, I want to identify intrinsic constraints, that is a model \mathcal{M} , that are coupled with the extrinsic constraints Σ according to some natural optimality conditions. In order to do so, let us consider the map that assigns to each sensorimotor loop its corresponding behaviour, the *behaviour map*:

$$\psi_{\sigma} : \Gamma \longrightarrow \Delta_{\mathcal{W}^{\infty} \times \mathcal{C}}, \quad \gamma \longmapsto \mathbb{P}(w^0, c^0; dw^1, dw^2, dw^3, \dots), \tag{5}$$

where we use the formula (2) to define the Markov kernel $\psi_{\sigma}(\gamma)$ by letting T go to ∞ . Now, say that we have given for all $\sigma \in \Sigma$ a set \mathcal{O}_{σ} of optimal or desired behaviours, such as behaviours with maximal predictive information [25] or maximal expected reward [8]. In this article, optimality of behaviours is not further specified and should not be confused with the optimality of a model, which plays the central role in this paper. Optimal models should, at least, satisfy the following natural sufficiency condition.

Fig. 4 Changes of the intrinsic mechanisms π and φ as result of changes of parameters such as synaptic coupling strengths



Definition 1 We say that a model \mathcal{M} is (geometrically) sufficient, if for all $\sigma \in \Sigma$ and all corresponding behaviours $\delta \in \mathcal{O}_\sigma$, there exists $\gamma \in \overline{\mathcal{M}}$ that generates that behaviour, that is $\psi_\sigma(\gamma) = \delta$.

Here, the bar over the set \mathcal{M} denotes its topological closure in Γ . In principle, this will depend on the underlying topology of Γ for which one has various natural choices. I am not going to address these topological questions in further detail. The main results of the next section refer to the case where the state sets are finite and therefore the topology is simply the standard one of a finite-dimensional real vector space. If the closure of the model \mathcal{M} equals all theoretically possible intrinsic mechanisms, that is $\overline{\mathcal{M}} = \Gamma$, then we say that the model is a universal approximator. This corresponds to the most flexible brain architecture which I already mentioned above [see Eq. (3)]. I argue that such a brain is not required for embodied agents in order to be universal at the behavioural level. There are sufficient models, in the sense of Definition 1, that are behaviourally equivalent to, but less complex than, a universal approximator. Clearly one has to specify the notion of complexity here. In any case, the general study of sufficient models in relation to their complexity provides one way to formally address the subject of cheap design which plays a central role within the field of embodied intelligence [19].

In order to be more precise, we consider as optimal behaviour sets \mathcal{O}_σ all behaviours that can in principle be generated in the context of the extrinsic constraints given by σ . Stated differently, we set \mathcal{O}_σ to be the image of ψ_σ and obtain the following sufficiency property of a model \mathcal{M} , which we then call an embodied universal approximator (for Σ):

$$\psi_\sigma(\overline{\mathcal{M}}) = \psi_\sigma(\Gamma) \quad \text{for all } \sigma \in \Sigma. \tag{6}$$

In general, an embodied universal approximator does not have to be concise. For instance, any universal approximator is in particular an embodied universal approximator. In order to be consistent with the idea of cheap design, we have to specify what we mean by “cheap” in terms of a complexity measure and then compare various embodied universal approximators with respect to their complexity. This will be exemplified in the next section.

4 Cheap Embodied Universal Approximation

The intention of this section is to exemplify the developed concepts and tools in terms of very specific theoretical results on cheap design which are closely related to the work [16]. In particular, I assume only finite sets \mathcal{W} , \mathcal{S} , \mathcal{C} , and \mathcal{A} , and, furthermore, I concentrate on the definition of policy models, ignoring other intrinsic constraints. To be more precise, consider the set

$$\Gamma = \Delta_{\mathcal{A}}^{\mathcal{C}} \times \Delta_{\mathcal{C}}^{\mathcal{C} \times \mathcal{S}}$$

of all possible intrinsic mechanisms π and φ [see Eq. (3)]. In general, a model \mathcal{M} defines a reduction of Γ in both components, that is $\mathcal{M} \subseteq \Gamma$. In this section, however, we study only reductions in the first component in terms of policy models $\mathcal{M}_{\text{policy}} \subseteq \Delta_{\mathcal{A}}^{\mathcal{C}}$ which correspond to joint controller models

$$\mathcal{M} = \mathcal{M}_{\text{policy}} \times \Delta_{\mathcal{C}}^{\mathcal{C} \times \mathcal{S}}. \tag{7}$$

We say that a policy model is an embodied universal approximator (for Σ), if the corresponding joint model (7)

has this property. Note that the restriction to policy models excludes the possibility of coupling the two intrinsic mechanisms π and φ . Such a coupling might provide a further way of reducing the complexity of the controller model.

4.1 General Selection of Policy Models

Consider the one-step kernel $\mathbb{P}(w, c; dw')$, which appears in (1). For finite sets, this reduces to

$$\mathbb{P}^\pi(w, c; w') = \sum_a \alpha(w, a; w') \pi(c; a), \quad c \in \mathcal{C}, w, w' \in \mathcal{W}, \quad (8)$$

where we indicate the dependence on π explicitly. We leave all other mechanisms of the sensorimotor loop fixed and consider two policies π_1 and π_2 that satisfy

$$\mathbb{P}^{\pi_1}(w, c; w') = \mathbb{P}^{\pi_2}(w, c; w') \quad \text{for all } c \in \mathcal{C} \text{ and all } w, w' \in \mathcal{W}. \quad (9)$$

It follows directly from (1) that the corresponding behaviours, defined by (5), are identical. In other words, if for all controller states c the distributions $\pi_1(c; \cdot)$ and $\pi_2(c; \cdot)$ give the same expectation values of the functions

$$\alpha_{w, w'} : a \mapsto \alpha_{w, w'}(a) := \alpha(w, a; w'), \quad w, w' \in \mathcal{W}, \quad (10)$$

then π_1 and π_2 will generate the same behaviour, assuming that all the other mechanisms are fixed. This redundancy allows us to construct policy models with smaller dimensionality without reducing the behavioural repertoire of the agent. In order to do so, we identify two distributions μ_1 and μ_2 on \mathcal{A} if they give the same expectation values of the functions $\alpha_{w, w'}, w, w' \in \mathcal{W}$, and thereby obtain a partition of $\Delta_{\mathcal{A}}$ into convex equivalence classes $[\mu], \mu \in \Delta_{\mathcal{A}}$. Given a policy π , we can now consider the family of classes $[\pi(c; \cdot)], c \in \mathcal{C}$, and any further π' satisfying $\pi'(c; \cdot) \in [\pi(c; \cdot)]$ will generate the same behaviour as π . Following the idea of cheap design, it is natural to define a model by choosing π' to be the one with the least complexity. Assuming that our complexity measure reflects the structure of a policy, and, furthermore, that structure corresponds to low entropy, we have the natural choice provided by the maximum entropy principle. More precisely, in each class $[\pi(c; \cdot)]$ we choose the distribution with the highest entropy $-\sum_{a \in \mathcal{A}} \pi(c; a) \ln \pi(c; a)$. This policy is uniquely determined, as the entropy is strictly concave and the classes are convex. We obtain our first parametrisation $\eta : \mathbb{R}^{\mathcal{C} \times \mathcal{W}^2} \rightarrow \Delta_{\mathcal{A}}^{\mathcal{C}}$,

$$\theta \mapsto \pi_\theta(c; a) := \frac{\exp\left(\sum_{w, w'} \theta_{w, w'}(c) \alpha_{w, w'}(a)\right)}{\sum_{a'} \exp\left(\sum_{w, w'} \theta_{w, w'}(c) \alpha_{w, w'}(a')\right)}. \quad (11)$$

This is a quite prominent model in statistics and information geometry, which is called *exponential family* [1]. Using general information-geometric arguments, it is obvious that the parametrisation (11) defines an embodied universal approximator for

$$\Sigma_{\mathcal{A}} := \{(\alpha, \beta) : \beta \in \Delta_{\mathcal{A}}^{\mathcal{W}}\}. \quad (12)$$

On the other hand, in most interesting cases the number $|\mathcal{C}| |\mathcal{W}|^2$ of parameters of this model will be a huge. It turns out, however, that this “full” representation of policies is highly redundant. One obvious redundancy follows from the fact that the $\alpha(w, a; w')$ are probability distributions in w' , that is $\sum_{w'} \alpha_{w, w'}(a) = 1$ for all $w \in \mathcal{W}$ and all $a \in \mathcal{A}$. As the space of constant functions on \mathcal{A} cancels out for each $c \in \mathcal{C}$, the number of parameters can be reduced to $|\mathcal{C}| (|\mathcal{W}|^2 - 1)$. I argue that a further, conceptually deeper, reduction can be expected in embodied systems. Formally, it is reflected by the fact that the sum

$$\sum_{w, w'} \theta_{w, w'}(c) \alpha_{w, w'}(a), \quad (13)$$

which appears in (11), can be strongly simplified due to linear dependence of the vectors $\alpha_{w, w'} \in \mathbb{R}^{\mathcal{A}}, w, w' \in \mathcal{W}$. Let us consider two obvious reasons for this.

1. Assume that the world state w' can not be reached from the world state w in one step, that is

$$\alpha(w, a; w') = 0 \quad \text{for all } a \in \mathcal{A}. \quad (14)$$

In this case, the corresponding term disappears from the sum (13). In the situation of an embodied agent, the majority of pairs (w, w') has this property, because the physical constraints of the sensorimotor loop exclude most of the transitions from w to w' in one step.

2. It is not necessary that (14) holds in order to ignore a term from the sum (13). It is already sufficient that there is a constant $r \in \mathbb{R}$ such that

$$\alpha(w, a; w') = r \quad \text{for all } a \in \mathcal{A}. \quad (15)$$

Although formally this property is a simple extension of (14), it highlights another important aspect. If $\alpha(w, a; w') = r > 0$ for all $a \in \mathcal{A}$ then w' can be reached from w in one step but in a way that does not involve the actuators. The transition from w to w' is not sensitive to the actuators, and therefore its representation does not play any role in our parametrisation (11).

Clearly, we can adjust the parametrisation (11) to the dimension $d_{\mathcal{A}}$ of the vector space

$$V_{\mathcal{A}} := \text{span}\{\alpha_{w, w'} : w, w' \in \mathcal{W}\} \subseteq \mathbb{R}^{\mathcal{A}},$$

which is at most $|\mathcal{A}|$. Furthermore, note that the constant functions are contained in $V_{\mathcal{A}}$. Therefore, there are $d_{\mathcal{A}} - 1$

functions $\alpha_k \in V_\alpha$, which we refer to as *feature vectors*, such that every policy of the structure (11) can be expressed in terms of a linear combination of the α_k and a constant function. This leads to the following simplification of the parametrisation (11).

Proposition 1 *The parametrisation $\eta : \mathbb{R}^{\mathcal{C}} \times \mathbb{R}^{d_\alpha-1} \rightarrow \Delta_{\mathcal{A}}^{\mathcal{C}}$,*

$$\theta \mapsto \pi_\theta(c; a) := \frac{\exp\left(\sum_{k=1}^{d_\alpha-1} \theta_k(c) \alpha_k(a)\right)}{\sum_{a'} \exp\left(\sum_{k=1}^{d_\alpha-1} \theta_k(c) \alpha_k(a')\right)}, \quad (16)$$

defines an embodied universal approximator for Σ_α .

In this parametrisation of the policies, we have at most $|\mathcal{C}|(d_\alpha - 1)$ parameters which can be much smaller than the number of parameters in the parametrisation (11). However, clearly they define the same policy model, which is a maximum entropy model. This was motivated by the idea that the maximum entropy distributions correspond to the least complex or structured ones. However, high entropy distributions tend to have a large support. We can take a different standpoint and ask the following question: If an agent can generate a behaviour with only a few actuator states, why should it involve more actuator states than that? In other words, if we consider actuator states to be costly then we should design the model in a way that the policies have only a minimal required number of actuator states. Interpreting, for simplicity of the argument, entropy as a cardinality measure, then we would rather prefer distributions with minimal entropy. From this perspective, the minimum entropy distributions appear more natural. Note, that, in contrast to the uniqueness of the maximum entropy distribution in a class $[\mu]$ where $\mu \in \Delta_{\mathcal{A}}$, there are many distributions that locally minimise the entropy in that class. On the other hand, it turns out that the cardinality of the support of all these distributions is upper bounded by d_α (see corresponding derivations in [2]). More precisely, the following holds: If for a policy π , all distributions $\pi(c; \cdot)$ are local minimisers of the entropy in the respective classes $[\pi(c; \cdot)]$, then

$$|\{a \in \mathcal{A} : \pi(c; a) > 0\}| \leq d_\alpha \quad \text{for all } c \in \mathcal{C}. \quad (17)$$

Clearly, a policy model $\mathcal{M}_{\text{policy}}$ is an embodied universal approximator, if it contains all policies that satisfy (17) in its closure. We now explicitly define such a policy model with

$$\dim \mathcal{M}_{\text{policy}} \leq 2|\mathcal{C}|d_\alpha.$$

Proposition 2 *For any injective function $f : \mathcal{A} \rightarrow \mathbb{R}$, for which $f(a) = f(a')$ only if $a = a'$, the parametrisation $\eta : \mathbb{R}^{\mathcal{C}} \times \mathbb{R}^{2d_\alpha} \rightarrow \Delta_{\mathcal{A}}^{\mathcal{C}}$,*

$$\theta \mapsto \pi_\theta(c; a) := \frac{\exp\left(\sum_{k=1}^{2d_\alpha} \theta_k(c) f^k(a)\right)}{\sum_{a'} \exp\left(\sum_{k=1}^{2d_\alpha} \theta_k(c) f^k(a')\right)}, \quad (18)$$

defines an embodied universal approximator for Σ_α (here, $f^k(a)$ denotes the k th power of $f(a)$).

The proof of this statement is based on [2, 14] and is given in the Appendix. The policy model defined by (18) has $2|\mathcal{C}|d_\alpha$ parameters, which is larger than the number $|\mathcal{C}|(d_\alpha - 1)$ of parameters of the previous model defined in Proposition 1. However, there are important differences. The feature vectors f^k , $k = 1, \dots, 2d_\alpha$, do not explicitly depend on α , but their number depends on the dimension of V_α . Therefore, they can be used for all mechanisms α' that have at most the same dimension $d_{\alpha'}$ as α . Furthermore, these feature vectors do not require much information to be specified. More precisely, the property of f that is required in Proposition 2 is a generic property of real functions on \mathcal{A} so that f can be chosen randomly. Furthermore, given such a function f as the first feature vector, all the other feature vectors are determined as the k -th powers of f . With the set

$$\Sigma_{d_\alpha} := \{(\alpha', \beta) : d_{\alpha'} \leq d_\alpha, \beta \in \Delta_{\mathcal{G}}^{\alpha'}\} \quad (19)$$

of extrinsic constrains, which is larger than Σ_α as defined by (12), we have the following direct implication of Proposition 2.

Corollary 1 *The policy model of Proposition 2 is an embodied universal approximator for Σ_{d_α} .*

4.2 From Models to Architectures

In the last section, we defined policy models for general unstructured sets \mathcal{C} and \mathcal{A} , and the policies were abstract objects without any reference to a mechanism that generates or computes an output $a \in \mathcal{A}$ given the input $c \in \mathcal{C}$. In this section, we extend our previous ideas to the setting of systems consisting of interacting nodes such as neuronal systems. However, we restrict attention to the actuators and assume, for simplicity, that we have $n \geq 1$ binary actuators with states -1 and $+1$, that is $\mathcal{A} = \{-1, +1\}^n$. For each non-empty subset $N \subseteq [n] := \{1, \dots, n\}$, the interaction among the actuators in N is described in terms of the monomial

$$a = (a_1, \dots, a_n) \mapsto \theta_N(c) \prod_{i \in N} a_i, \quad (20)$$

where $\theta_N(c)$ is the interaction coefficient or strength of the actuators in N , which we assume to be dependent on the controller state $c \in \mathcal{C}$. We refer to the cardinality of N as the *order* of the interaction. For $|N| = 2$, we have the important special case of a pairwise interaction, which is of particular interest within the field of neural networks. There, the interaction coefficients are usually interpreted as the synaptic connection strengths between the neurons. If we want to incorporate interactions among nodes of all

subsets $N \subseteq [n]$, then we have to consider the following sum of monomials (20):

$$\sum_{\emptyset \neq N \subseteq [n]} \theta_N(c) \prod_{i \in N} a_i. \tag{21}$$

Note that, for each $c \in \mathcal{C}$ we have $2^n - 1$ parameters $\theta_N(c)$, so that the overall number of parameters is $|\mathcal{C}|(|\mathcal{A}| - 1)$, which coincides with the dimension of the polytope $\Delta_{\mathcal{A}}^{\mathcal{C}}$. Clearly, any function $\mathcal{A} = \{-1, +1\}^n \rightarrow \mathbb{R}$ can be written in the form (21). By setting in (21) all $\theta_N(c)$ to zero, if $|N| > k$, we reduce this space of functions to the space of functions that involve only interaction of at most order k . This defines the following policy model, which we refer to as k -interaction model:

$$\theta \mapsto \pi_{\theta}(c; a) = \frac{\exp\left(\sum_{\substack{\emptyset \neq N \subseteq [n] \\ |N| \leq k}} \theta_N(c) \prod_{i \in N} a_i\right)}{\sum_{a'} \exp\left(\sum_{\substack{\emptyset \neq N \subseteq [n] \\ |N| \leq k}} \theta_N(c) \prod_{i \in N} a'_i\right)}. \tag{22}$$

Given this increasing family of models, ordered according to k , we now address the problem of finding the sufficient order of interaction so that the corresponding k -interaction model is an embodied universal approximator for Σ_{d_x} . For a first simple estimate, we consider again the policy model of Proposition 2. Given that we now assume $\mathcal{A} = \{-1, +1\}^n$, it is possible to define the function f as a linear function, that is $f(a) = f(a_1, \dots, a_n) = \sum_{i=1}^n w_i a_i$ so that $f(a) \neq f(a')$ whenever $a \neq a'$. Note that this follows from the finiteness of \mathcal{A} . As a linear function defined on \mathbb{R}^n , f is never injective, except for $n = 1$. For the k th powers of f we obtain

$$\pi_{\theta}(c; h, a) := \frac{\exp\left(\sum_{\substack{1 \leq i \leq m \\ 1 \leq j \leq n}} \theta_{i,j}(c) h_i a_j + \sum_{1 \leq i \leq m} \theta_i(c) h_i + \sum_{1 \leq j \leq n} \theta_j(c) a_j\right)}{\sum_{h', a'} \exp\left(\sum_{\substack{1 \leq i \leq m \\ 1 \leq j \leq n}} \theta_{i,j}(c) h'_i a'_j + \sum_{1 \leq i \leq m} \theta_i(c) h'_i + \sum_{1 \leq j \leq n} \theta_j(c) a'_j\right)}. \tag{23}$$

$$f^k(a_1, \dots, a_n) = \left(\sum_{i=1}^n w_i a_i\right)^k = \sum_{i_1, \dots, i_k} w_{i_1} \cdots w_{i_k} a_{i_1} \cdots a_{i_k}.$$

This implies

$$\begin{aligned} \sum_{k=1}^{2d_x} \theta_k(c) f^k(a_1, \dots, a_n) &= \sum_{k=1}^{2d_x} \theta_k(c) \sum_{i_1, \dots, i_k} w_{i_1} \cdots w_{i_k} a_{i_1} \cdots a_{i_k} \\ &= \sum_{\substack{N \subseteq [n] \\ |N| \leq 2d_x}} \theta_N(c) \prod_{i \in N} a_i \end{aligned}$$

with corresponding coefficients $\theta_N(c)$. This proves that the policy model of Proposition 2 is contained in the $(2d_x)$ -interaction model, and that therefore the latter is also an embodied universal approximator. However, it turns out that this can be considerably improved. It directly follows from a result by Kahle [12] that $2d_x$ can be replaced by a much smaller interaction order, which defines the following policy model.

Proposition 3 *With*

$$k(\alpha) := \lceil \log_2(d_x + 1) \rceil,$$

the $k(\alpha)$ -interaction model is an embodied universal approximator for Σ_{d_x} ($\lceil x \rceil$ denotes the smallest integer $\geq x$).

In neural networks theory, the interaction order is generally considered to be pairwise. This is valid for the policy model of Proposition 3 only in the trivial case $d_x = 1$. If we aim at a neuronal interpretation of policy models then we have to represent higher order interactions in terms of pairwise interactions, which is always possible by extending the system appropriately. One important special case in line with this general idea is given by the so-called *restricted Boltzmann machine* (RBM). In order to define it, we extend the above system of n nodes by further m so-called *hidden nodes* with the same state space $\{-1, +1\}$. The overall state is then given as a pair $(h, a) = (h_1, \dots, h_m, a_1, \dots, a_n) \in \{-1, +1\}^{m+n}$. We now consider the following family of kernels, which involves at most pairwise interactions between the hidden nodes and the actuators:

Finally, we marginalise out the hidden nodes and obtain the following policy model, which is, for each $c \in \mathcal{C}$, a restricted Boltzmann machine:

$$\theta \mapsto \pi_{\theta}(c; a) = \sum_h \pi_{\theta}(c; h, a). \tag{24}$$

Therefore, we call this policy model *controlled restricted Boltzmann machine* (see Fig. 5). Due to general results by Le Roux and Bengio [22] and Montúfar and Ay [15], it is known that any distribution of support cardinality smaller than or equal to d_x can be represented by an RBM with

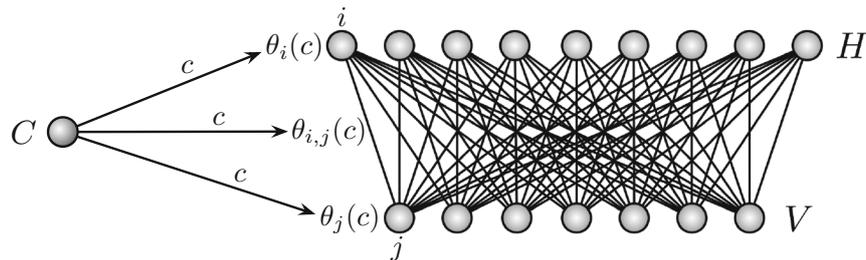


Fig. 5 The architecture of a controlled restricted Boltzmann machine. It consists of a restricted Boltzmann machine with the set $V = \{1, \dots, n\}$ of visible nodes and the set $H = \{1, \dots, m\}$ of hidden nodes. The interaction is pairwise, where each node in V interacts with each node in H . No lateral interactions within V and within H are assumed. The interaction coefficients θ_{ij} , θ_i , and θ_j , where $i \in H$ and $j \in V$, are all functions of the controller state c

$d_x - 1$ hidden nodes. Together with (17), this directly implies the following result.

Proposition 4 A controlled restricted Boltzmann machine, defined by (23) and (24), with $m = d_x - 1$ hidden nodes is an embodied universal approximator for Σ_{d_x} .

The dependence of the parameters $\theta_{i,j}(c)$, $\theta_i(c)$, and $\theta_j(c)$ on the controller state c can be quite complicated. Assuming that the controller also has a composite structure with k binary nodes, it is possible to represent this dependence in terms of pairwise interactions between the controller nodes and the hidden nodes. This leads to the definition of a *conditional restricted Boltzmann machine* [17]. Cheap control with such machines has been theoretically and experimentally studied in [16]. This study is based on the notion of embodiment dimension, which is a refinement of the dimension d_x used in the present article.

5 Conclusions

In order to approach a general theory of embodied agents, I have introduced a formal model of the sensorimotor loop, which specifies its intrinsic and extrinsic mechanisms, building on previous work [7, 13]. The extrinsic mechanisms represent the embodiment constraints of the system which can be utilised by appropriate adjustment of the intrinsic mechanisms in order to express useful behaviours. This requires some degree of flexibility of the control architecture, which I formalised in terms of a sufficiency notion. As a particular case, I studied in more detail embodied universal approximation and the corresponding design of controller models in terms of geometry. However, I argued that sufficiency should not be the only requirement involved in systems design. In order to address the notion of cheap design within the field of embodied intelligence, we have to identify sufficient controller architectures with low complexity. Clearly, one would assume that selection

pressure has generated such architectures of naturally evolved brains, which can cope with limitations of mass and energy resources. For instance, the field of sparse coding has provided evidence for internal representations of external stimuli in terms of sparse neuronal activity [18]. However, there are many aspects that contribute to the overall complexity of a controller model, and we are far from a conclusive definition of complexity that would account for the right notion of cheap control. Therefore, I was not very precise in this regard and used a few complexity notions as motivation for the policy models that I have defined. In this context, low complexity of a policy meant: 1. high entropy of the policy, 2. low number of actuator states used by the policy, 3. low interaction order among the actuators, and 4. low number of hidden nodes in a restricted Boltzmann machine.

There is a further important aspect of complexity, which I did not address explicitly. If we design policy models, we have to distinguish between the complexity of the model itself and the complexity of a particular policy, taken from that model. One can design models for which the individual policies have low complexity, but the overall model is hard to describe. However, it is problematic to implement such a structure. In nature, its information has to be transmitted through genetic inheritance. This information transmission clearly has a limited capacity, which is, for instance, modulated by the mutation rate. Therefore, robustness issues have also to be taken into account [5, 20]. To conclude, the right choice of a model should balance the complexities of both, the model itself and the policies that are implemented by the model. This is a quite natural idea within complexity theory [6, 10, 11, 21, 23].

Acknowledgments I would like to thank Keyan Ghazi-Zahedi, Guido Montúfar, and Johannes Rauh for many stimulating discussions on the subject of embodied intelligence, and, in particular, systems design. The proof of Proposition 2, presented in the Appendix, uses an argument by Johannes Rauh, which was not used in the original proof [14], leading to an improvement of the original result. This work has been supported by the DFG Priority Program *Autonomous Learning*.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Appendix

Proof of Proposition 2 We consider a policy π with the property that the supports \mathcal{A}_c of its distributions $\pi(c; \cdot)$, $c \in \mathcal{C}$, on \mathcal{A} have at most d_x elements. For each $c \in \mathcal{C}$, we define the function $g_c : \mathcal{A} \rightarrow \mathbb{R}$,

$$a \mapsto g_c(a) := \prod_{a' \in \mathcal{A}_c} (f(a) - f(a'))^2 = \sum_{k=0}^{2d_x} \theta_k^{(0)}(c) f^k(a).$$

Obviously, $g_c \geq 0$, and $g_c(a) = 0$ if and only if $a \in \mathcal{A}_c$. Furthermore, we consider the polynomial $h_c : \mathbb{R} \rightarrow \mathbb{R}$, $t \mapsto \sum_{k=0}^{d_x-1} \theta_k^{(1)}(c) t^k$, that satisfies

$$\sum_{k=0}^{d_x-1} \theta_k^{(1)}(c) f^k(a) = h_c(f(a)) = \ln \pi(c; a), \quad a \in \mathcal{A}_c.$$

Obviously, the one parameter family

$$\pi_\lambda(c; a) = \frac{e^{h_c(f(a)) - \lambda g_c(a)}}{\sum_{a' \in \mathcal{A}_c} e^{h_c(f(a')) - \lambda g_c(a')}}.$$

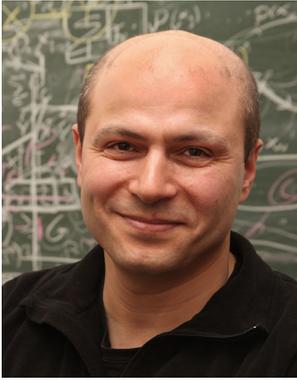
of policies is contained in the model defined by (18). Note that we can ignore the term for $k = 0$ because this gives a constant which cancels out. Furthermore, it is easy to see that

$$\lim_{\lambda \rightarrow \infty} \pi_\lambda(c; a) = \lim_{\lambda \rightarrow \infty} \frac{e^{h_c(f(a))} (e^{-g_c(a)})^\lambda}{\sum_{a' \in \mathcal{A}_c} e^{h_c(f(a'))} (e^{-g_c(a')})^\lambda} = \pi(c; a). \quad (25)$$

This proves that π is in the closure of the model defined by (18). \square

References

- Amari S, Nagaoka H (2000) Methods of information geometry. American Mathematical Society, Oxford University Press
- Ay N (2002) An information-geometric approach to a theory of pragmatic structuring. *Ann Probab* 30(1):416–436
- Ay N, Jost J, Lê HV, Schwachhöfer L (2015) Information Geometry, Springer (submitted)
- Ay N, Knauf A (2007) Maximizing multi-information. *Kybernetika* 42(5):517–538
- Ay N, Krakauer DC (2007) Geometric robustness theory and biological networks. *Theory Biosci* 125(2):93–121
- Ay N, Müller M, Szkoła A (2010) Effective complexity and its relation to logical depth. *IEEE Trans Inf Theory* 56(9):4593–4607
- Ay N, Zahedi K (2014) On the causal structure of the sensorimotor loop. In: Prokopenko M (ed) Guided self-organization: inception. Springer, Berlin, Heidelberg
- Ay N, Montúfar G, Rauh J (2012) Selection criteria for neuro-manifolds of stochastic dynamics. Springer, Post-conference proceedings Advances in Cognitive Neurodynamics (III)
- Brooks RA (1991) Intelligence without representation. *Artif Intell* 47(1–3):139–159
- Gell-Mann M, Lloyd S (1996) Information measures, effective complexity, and total information. *Complexity* 2:44–52
- Jost J (2004) External and internal complexity of complex adaptive systems. *Theory Biosci* 123:69–88
- Kahle T (2010) Neighborliness of marginal polytopes. *Contrib Algebra Geom* 51(1):45–56
- Klyubin AS, Polani D, Nehaniv CL (2004) Tracking information flow through the environment: Simple cases of stigmerg. In: Pollack J (ed) Artificial Life IX: Proceedings of the Ninth International Conference on the simulation and synthesis of living systems, pages 563–568. MIT Press
- Matúš F, Ay N (2003) On maximization of the information divergence from an exponential family. In: Vejnarová J (ed) Proceedings of WUPES'03, University of Economics Prague pp 199–204
- Montúfar G, Ay N (2011) Refinements of universal approximation results for deep belief networks and restricted Boltzmann machines. *Neural Comput* 23(5):1306–1319
- Montúfar G, Zahedi K, Ay N (2015) A theory of cheap control in embodied systems. *PLOS Comput Biol*. [arXiv:1407.6836](https://doi.org/10.1371/journal.pcbi.1004636) (in press)
- Montúfar G, Ay N, Zahedi K (2015) Geometry and expressive power of conditional restricted Boltzmann machines. *J Mach Learn Res*. [arXiv:1402.3346](https://arxiv.org/abs/1402.3346) (in press)
- Olshausen BA, Field DJ (1996) Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381:607–609
- Pfeifer R, Bongard JC (2006) How the body shapes the way we think: a new view of intelligence. The MIT Press (Bradford Books), Cambridge
- Rauh J, Ay N (2013) Robustness, canalising functions, and systems design. *Theory in Biosci*. doi:10.1007/s12064-013-0186-3
- Rissanen J (1989) Stochastic complexity in statistical inquiry. World Scientific, Singapore
- Le Roux N, Bengio Y (2008) Representational power of restricted Boltzmann machines and deep belief networks. *Neural Comput* 20(6):1631–1649
- Vapnik VN (1998) Statistical learning theory. Wiley, New York
- Zahedi K, Ay N (2013) Quantifying morphological computation. *Entropy* 15(5):1887–1915. doi:10.3390/e15051887
- Zahedi K, Ay N, Der R (2010) Higher coordination with less control: a result of information maximisation in the sensorimotor loop. *Adapt Behav* 18:338–355



Short CV of Nihat Ay I studied mathematics and physics at the Ruhr University Bochum and received my Ph.D. in mathematics from the University of Leipzig in 2001. In 2003 and 2004, I was a postdoctoral fellow at the Santa Fe Institute and at the Redwood Neuroscience Institute (now the Redwood Center for Theoretical Neuroscience at UC Berkeley). After my postdoctoral stay in the USA I became a faculty member at the Friedrich Alexander

University in Erlangen where I earned my habilitation in

mathematics. Since September 2005, I have worked at the Max Planck Institute for Mathematics in the Sciences in Leipzig where I am heading the research group Information Theory of Cognitive Systems. I am also affiliated with the University of Leipzig as Honorary Professor and with the Santa Fe Institute as Part-time Resident Professor.