IOP Conference Series:
Materials Science and
Engineering

An Information-Theoretic Perspective on Acting Agents

To cite this article: Carlotta Langer and Nihat Ay 2024 *IOP Conf. Ser.: Mater. Sci. Eng.* **1321** 012008

An Information-Theoretic Perspective on Acting Agents

Carlotta Langer[1] **and Nihat Ay**[1,2,3]

[1] Hamburg University of Technology, Hamburg, Germany
[2] Santa Fe Institute, Santa Fe, USA
[3] Leipzig University, Leipzig, Germany

E-mail: `carlotta.langer@tuhh.de`

Abstract. Every embodied intelligent agent constantly interacts in its own way with its environment. It perceives information about the world, processes this information in its controller, e.g. its brain, and sends signals to the body, which then interacts with its surroundings. Here we give a thorough introduction to an information-theoretic approach that aims at quantifying these information flows among an agent's body, brain and environment. Thereby we measure how much an information flow impacts the behavior of an agent. For instance, we might ask the question whether the brain is essential for a certain behavior or whether it is a reflexive stimuli response. In order to measure these interactions we model the system using the sensorimotor loop and assess the different information flows by information-theoretic measures. This includes measures related to Morphological Computation and the complexity of the controller. Morphological Computation describes the reduction of computational cost for the controller that results from the interaction with the environment. We include results from an implementation of this framework with simplistic simulated robots. Experiments with these robots show, for instance, that an agent with a well-adapted body does not rely as much on its controller. Furthermore, learning a new task requires an increased information integration in the controller.

Keywords: Information Theory, Morphological Computation, Integrated Information

1. Introduction

Embodied agents, whether they are biological or artificial, constantly interact with their environments. In Figure 1 a sketch of such an agent is depicted. There the agent senses information about its environment through its eyes and processes this information in its controller, which is the brain in case of biological agents and could be an artificial neural network in case of a robot. The information from the sensors can then lead to a direct stimuli response, a reflexive action such as a movement of a hand, or the controller can send commands to the hands, which in turn influence the environment.

When an intelligent embodied agent, like an animal, is faced with a changed environment it has to adapt and learn new behaviors. It has to figure out how to best use the properties of its own body to cope with the new surroundings. There are various theories that aim to explain how exactly animals learn. Here we take a step back from this problem and do not ask "how" an agent learns, but instead which parts

IOP Conf. Series: Materials Science and Engineering **1321** (2024) 012008 doi:10.1088/1757-899X/1321/1/012008

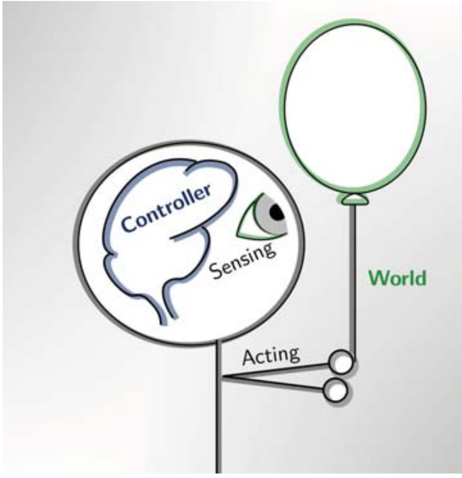

Figure 1. Sketch of an embodied agent that interacts with its environment.

of an agent are most important for its behavior. If the agent is catching an object, for example, is this behavior mostly governed by the commands from the brain of the agent or is this a reactive stimuli response? Does this change during the process of learning this behavior?

In our daily experiences we have an intuition of how important the influences from the brain or our body are. Consider a skilled musician playing an instrument. The musician does not need to actively think about where to place each finger on the instrument, since the movements that translate to the desired melody have already been learned and can be applied almost automatically. On the other hand, someone who is new to an instrument needs to think carefully about every movement, hand placement and learn through repetitions. So, our hypothesis would be that an agent that learns a new task needs first a high involvement of the brain, which then gradually decreases.

Now assume we have a robot learning to play the same instrument. There we are able to directly observe the signals processed in the controller, e.g. a neural network, and measure them using information theory. In order to learn we would expect it to go through similar stages of development. This could be an indication of how natural the learning process of an artificial agent is or in which way it differs from our expectations. In [34] we discuss a simplistic setting with learning simulated robots and there we observe the dynamics hypothesized above. There we conclude that an agent first needs a high involvement from a complex controller to understand its environment and the more it can exploit the interactions between body and environment, the less controller complexity is needed.

In this article we give a thorough introduction to the information-theoretic methodology for tracing information flows within sensorimotor interactions, thereby following the Ph.D. thesis of one of the authors [32]. In the following sections we therefore first give a quick introduction to information theory and then we define the sensorimotor loop in the Section 3. The sensorimotor loop models the different parts

IOP Conf. Series: Materials Science and Engineering **1321** (2024) 012008 doi:10.1088/1757-899X/1321/1/012008

of an acting agent in its environment. Afterwards, in Section 4, we introduce a way to measure the importance of different information flows. Hereby, we consider pair-wise as well as synergistic influences. Two interesting concepts, that can be measured in this way, are Morphological Computation and Integrated Information, introduced in Section 5 and 6, respectively. The former describes the reduction of computational cost for the controller that results from the interaction of the body with its environment and the latter aims to quantify consciousness by a system's ability to integrate information. In Section 7 we review previous results that we gained by applying this framework to simple simulated agents. We conclude with a discussion in Section 8.

2. Information Theory

In 1948 C. E. Shannon published the article "A Mathematical Theory of Communication" [51], developed further into a book with W. Warren as "The Mathematical Theory of Communication" [52], which provides the foundation of information theory. We give a short introduction in this section, a more thorough treatment of information theory can be found in, for example, [15, 35].

The formalization of a communication system is sketched in Figure 2. A message gets passed from an information source through a transmitter, then the signal might be disturbed by noise before it reaches the receiver and then its destination.

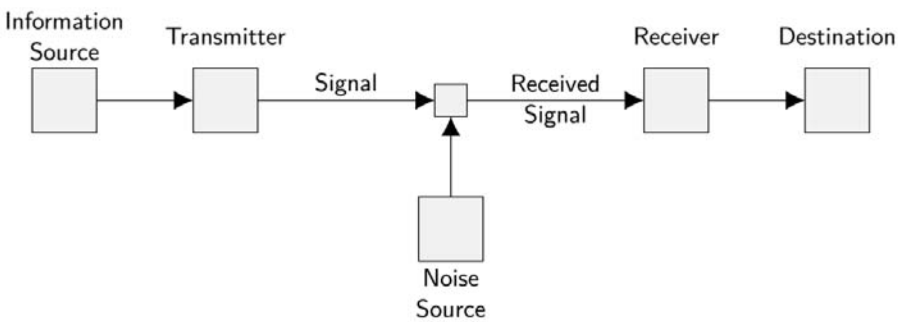

Figure 2. Reproduction of Shannon's general communication system, Figure 1 in [51].

Within his theory Shannon defines the optimal rate of data transportation and the maximal compression without loss of information in order to improve the transmission of information. One important concept in this context is the entropy.

Definition 2.1 (Entropy)**.** The entropy H of a discrete random variable X with the state space $\mathcal{X}$ and the probability mass function $p(X)$ is defined as follows:

$$H(X) = -\sum_{x \in \mathcal{X}} p(x) \log p(x)$$

with the convention $0 \log 0 = 0$.

The entropy is non-negative and is upper bounded by the entropy of the uniform distribution over $\mathcal{X}$. The base of the logarithm merely determines in what units we measure information, bits for the base 2 or nats for the base e, and this has no further implication on the properties of H.

The entropy is a lower bound for the compression and it can be thought of as the average reduction of uncertainty from observing an outcome x, when we initially know the probabilities $p(x)$. If we have a random variable with a uniform distribution, for example given by a fair die, then observing each side of the die is equally likely. Hence, the reduction of uncertainty would be large, roughly 2.585 bits. If, however, we have a weighted die that falls on one side 95% of all throws, then we would expect to observe this side and our average reduction of uncertainty is low, in this case roughly 0.402 bits.

In order to quantify the difference between two probability distributions we now define the Kullback-Leibler-divergence, also called KL-divergence or relative entropy. It measures how much the uncertainty of the random variable increases if we use q as probability distribution instead of the true distribution p.

Definition 2.2 (Relative Entropy)**.** The relative entropy, or KL-divergence D of two probability mass functions p and q with the state space $\mathcal{X}$ is defined as follows:

$$D(p \parallel q) = \sum_{x \in \mathcal{X}} p(x) \log \frac{p(x)}{q(x)}$$

with the conventions that $0 \log \frac{0}{q(x)} = 0$ for $q(x) \geq 0$ and $p(x) \log(\frac{p(x)}{0}) = \infty$ for $p(x) > 0$.

The KL-divergence is non-negative and equals zero if and only if the two arguments are equal, meaning that $p(x) = q(x)$, for all $x \in \mathcal{X}$. Note that the KL-divergence is not symmetric, hence this is not a metric.

Shannon defines the channel capacity by using the mutual information, which is the relative entropy between a distribution $p(x, y)$ and the product distribution $p(x)p(y)$ on $\mathcal{X} \times \mathcal{Y}$. The maximum of this quantity with respect to $p(x)$ leads to the maximal rate at which information can be sent through a communication channel, characterized by $p(y|x)$, with a vanishingly low probability of error.

In the following sections we will also use the conditional mutual information.

Definition 2.3 (Conditional Mutual Information)**.** Let (X_1, X_2, X_3) be a random vector on $\mathcal{X} = \mathcal{X}_1 \times \mathcal{X}_2 \times \mathcal{X}_3$ and let $p(x_1, x_3) > 0$, for all $(x_1, x_3) \in \mathcal{X}_1 \times \mathcal{X}_3$. The conditional mutual information of the random variables X_1 and X_2 given X_3 is defined as

$$I(X_1; X_2|X_3) = \sum_{(x_1,x_2,x_3) \in \mathcal{X}} p(x_1, x_2, x_3) \log \frac{p(x_1, x_2|x_3)}{p(x_1|x_3)}.$$

The conditional mutual information is 0 if and only if X_1 and X_2 are independent given X_3, this can be written as $X_1 \perp\!\!\!\perp X_2|X_3$.

The main theorems of classical information theory, like the noisy channel coding theorem, are asymptotic results, but the information-theoretic quantities above give also rise to many advanced theories. Here, we will only mention a few of them to give an intuition to the variety of applications of information theory.

IOP Conf. Series: Materials Science and Engineering **1321** (2024) 012008 doi:10.1088/1757-899X/1321/1/012008

In machine learning the evidence lower bound of the variational auto-encoder (VAE) is given by relative-entropy and cross-entropy terms [25]. In cybernetics, for instance, Ashby defines his law of requisite variety in terms of the Shannon entropy in [5]. He postulates that a system needs to have a greater than or equal number of states in its control mechanism compared to the system being controlled in order to be stable.

Furthermore, the well-known "active inference" framework aims to explain the behavior of sentient agents as them minimizing the free energy, which is defined using relative entropy terms. The free energy function is a variational bound on the surprisal, given by $-\log p(x)$, also called Shannon information, [42].

In Section 6 we additionally introduce the Integrated Information Theory, which aims at measuring the consciousness of a system.

3. Sensorimotor Loop

In order to analyze an acting agent in its environment we use the sensorimotor loop to model the interactions among an agent's sensors, actuators, its controller and the environment. This model is also called perception-action loop or action perception cycle [26, 27, 49, 53].

The agent receives information about the world through its sensors and can then process this information in its controller. It can also send signals from the sensors to the actuators, which corresponds to a direct stimuli response. The actuator can additionally receive commands from the controller and the signals from the internal actuators lead to actions and influence the environment of the agent. A sketch of the sensorimotor loop is given in Figure 3 on the left.

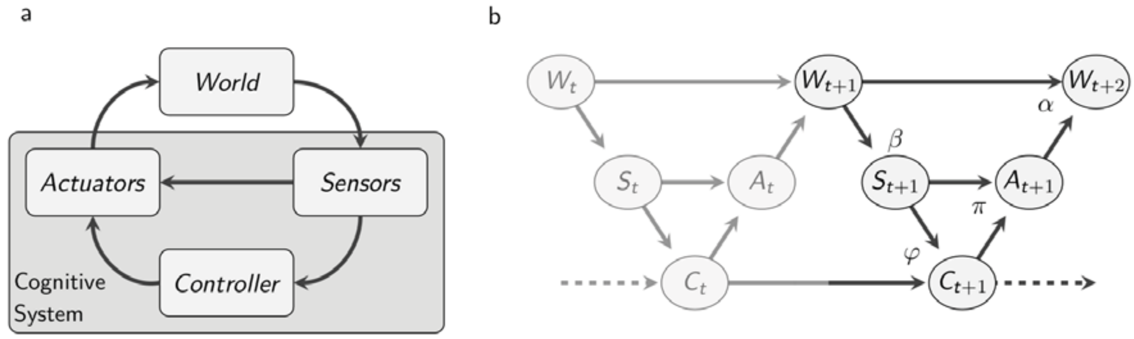

Figure 3. (a) The sensorimotor loop and (b) the graphical representation of the sensorimotor loop in discrete timesteps.

Here, we further model the system as a Bayesian network, similar to [12, 10, 28]. We assume discrete timesteps, which leads to the graphical representation on the right hand side of Figure 3. The different parts of the sensorimotor loop are modeled as discrete random variables and the interactions among them are given by the Markov kernels α, β, π and φ, as noted on the right of Figure 3. These Markov kernels are stochastic maps that determine the mechanisms of the variables. The kernel α, for instance, is called the world dynamics kernel in [18] and describes the probability of a state w_{t+2} given w_{t+1} and a_{t+1} by $\alpha(w_{t+2}|w_{t+1}, a_{t+1})$.

The whole system is then described by a discrete Markov Process $(X_t)_{t\in\mathbb{N}}$ with

$$X_t = (W_t, S_t, A_t, C_t)$$

and the state space $\mathcal{X} = \mathcal{W}\times\mathcal{S}\times\mathcal{A}\times\mathcal{C}$. The transition probability of the whole system from time step t to $t+1$ can be factorized as follows

$$p(x_{t+1}|x_t) = \alpha(w_{t+1}|w_t, a_t)\beta(s_{t+1}|w_{t+1})\varphi(c_{t+1}|c_t, s_{t+1})\pi(a_{t+1}|s_{t+1}, c_{t+1})$$

for all $(x_t, x_{t+1}) \in \mathcal{X}\times\mathcal{X}$. Using this description of the sensorimotor loop we can now define the measures of the different information flows.

4. Measuring an Information Flow

In the previous section we discuss the sensorimotor loop, a formalization of the information flows among the different parts of an agent and its environment. Next, we introduce a method to measure the strength of these information flows. The main idea our measures follow is that the strength of a connection is given as the minimal difference between a full system and the set of systems without this connection. We formalize this as follows.

Definition 4.1 (Measure $\Psi_{\mathcal{M}}$). Let p be a full distribution factoring according to the sensorimotor loop and let $\mathcal{M}$ be a set of positive probability distributions that factor according to the sensorimotor loop but are lacking the information flow that we want to measure. Then we define the measure $\Psi_{\mathcal{M}}$ by minimizing the KL-divergence between $\mathcal{M}$ and the full distribution

$$\Psi_{\mathcal{M}} = \inf_{q\in\mathcal{M}} D(p \parallel q).$$

This equips us with a technique to measure all the information flows among the different parts of the sensorimotor loop. Measures that use this approach to quantify an information flow are discussed in, for example, [22, 19, 20, 32]. In a similar setting, in [28], the amount of control an agent has over its environment is measured as "empowerment". The authors of [50] use such a measure to improve the efficiency of a reinforcement learning algorithm.

An alternative to this observational technique is to actively perturb the system, which would model an intervention and can be formalized by Pearl's *do*-calculus, described in [44, 43, 45]. This method is used in [11] to define causal information flows. Different approaches to quantify the causal strength of an information flow are discussed in [23].

As an example of our approach we now discuss the situation sketched in Figure 4. There we have a simplified agent without a controller and without sensors. In this case we have an agent that generates actions and influences its environment through them.

Consider a toy car with a broken remote that continuously moves forward and steers left and right with a fixed probability. Now, depending on the environment we let this car run in, the actions might have a large impact on the state of the environment, which includes the position of the car. In an open room the position of the car in the next point is determined by its last position and the selected actions.

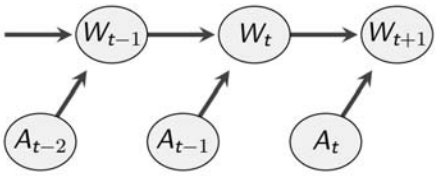

Figure 4. Acting agent influencing its environment.

If, however, this car is in a labyrinth with narrow passages, then the next position of the car is mostly determined by the environment. Now, we want to ask the question how to measure the strength of the influence of the world, the actuators or their synergistic influence. We start with quantifying a pair-wise connection.

First, we want to measure the influence of A_t on W_{t+1} given W_t and we define a set of distributions $\mathcal{M}$ that describe agents in which this connection does not exist, sketched as dashed arrow in Figure 5. This means that the kernel α only depends on W_{t+1} and W_t, hence $\alpha(w_{t+1}|a_t, w_t) = \alpha(w_{t+1}|w_t)$ for all $(w_{t+1}, w_t, a_t) \in \mathcal{W} \times \mathcal{W} \times \mathcal{A}$. The minimization in Definition 4.1 has the following closed form solution in this case

$$\Psi_{\mathcal{M}} = I(W_{t+1}; A_t|W_t).$$

The equality above can be easily proven by using that the cross-entropy is greater or equal to the entropy, as applied in a similar setting in Proposition 2 of [31].

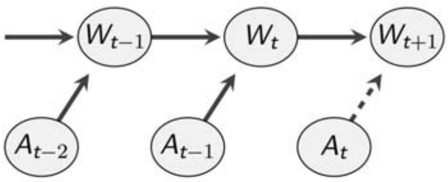

Figure 5. Acting agent for which the connection between A_t and W_{t+1} should be quantified.

Since $\Psi_{\mathcal{M}}$ can be written as a conditional mutual information we can conclude that this measure is 0 if and only if W_{t+1} is independent of A_t given W_t. In that case the action of the agent would have no influence on the environment at all.

In addition to the pair-wise interactions we are also able to define a measure for the synergistic influences based on triple-wise interactions, as sketched on the left of Figure 6. If the toy car steers into an obstacle and pushes it away, then this outcome is not determined by the action or the environment alone, but can only be explained by considering the information from both, W_t and A_t at the same time.

Now, in order to measure the synergistic influence of W_t and A_t on W_{t+1} we need to define $\mathcal{M}$ as a set of all those distributions in which there exist no three-way interactions

IOP Conf. Series: Materials Science and Engineering **1321** (2024) 012008 doi:10.1088/1757-899X/1321/1/012008

among the variables. These are sketched on the right of Figure 6. We are able to define these distributions as those for which there exist two non-negative functions f_1, f_2 such that

$$p(w_{t+1}|w_t, a_t) = f_1(w_{t+1}, a_t)f_2(w_{t+1}, w_t),$$

for all $(w_{t+1}, w_t, a_t) \in \mathcal{W} \times \mathcal{W} \times \mathcal{A}$.

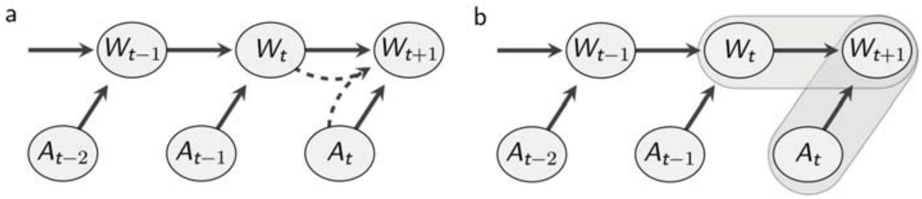

Figure 6. Sketch of an acting agent with a synergistic influence in (a) and one without such an influence in (b).

Minimizing the KL-divergence in Definition 4.1 does not lead to a closed form solution here. Instead, this can be calculated using, for instance, the iterative scaling algorithm, defined in [16]. This is discussed in more detail in the context of a synergistic measure for an internal prediction in [34] and a synergistic measure for Morphological Computation in [20]. Next we introduce the notion of Morphological Computation in more detail.

5. Morphological Computation

An embodied agent can interact with its environment in different ways depending on the properties of its body. When we grasp a fragile object, for instance, we can vary the amount of pressure that we are able to apply without damaging the object or loosing grip, because of the soft tissue of our skin. However, a robot with a metal hand needs to carefully calculate the amount of applied pressure. This intuitive concept, that the morphology of the body can alleviate the computational burden of the brain, is called Morphological Computation. There exist various definitions of Morphological Computation, especially the use of the term “computation” is debated, as discussed in more detail in, for example, [18, 38].

The importance of considering an embodied agent in a real environment in order to understand its behavior is stressed in [13, 14] and the concept of Morphological Computation was introduced in [48]. Remarkable cases of Morphological Computation can be seen in the field of soft robotics, for instance in [39] in which die authors present an octopus inspired robotic arm.

Additionally, Pfeifer and Bongard formulate in [46] the “cheap design principle”. It states that the body of an agent should be constructed with the goal to exploit the properties of their environment as best as possible. The authors of [37] use the information-theoretic setting discussed here in order to formalize the notion of cheap control architectures in the context of universal approximation.

In Section 7 we discuss results of Morphological Computation measures in experiments with simple simulated agents. Hence, we now introduce one class of measures for Morphological Computation that stem from the concepts introduced in the previous section. Figure 7 is a sketch of the three different information flows that we can measure and associate with Morphological Computation. The measure used in Section 7 quantifies the information flow through the world, given the action, marked with (1) in Figure 7. This is also discussed in [21, 18].

On the other hand, we can also measure the impact that the agent has on its environment, sketched as (2) in Figure 7. Subtracting the normalized strength of this connection from one leads in [21, 18] to the definition of another Morphological Computation measure.

The connection (3) in Figure 7, the interplay of the influences of W_t and A_t on W_{t+1}, is defined in [20] as a synergistic measure for Morphological Computation.

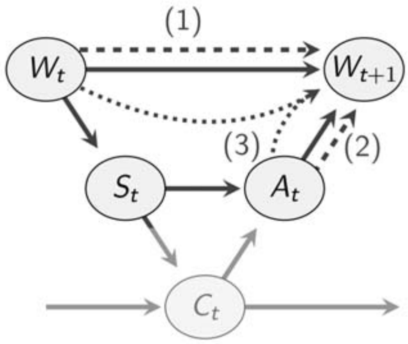

Figure 7. Different information flows that lead to a measure for Morphological Computation.

Since it is often not possible to access the world states directly, we can instead use only the information about the environment that is accessible to agents, given by the sensory states. The justification for using S instead of W is discussed in [33].

6. Controller Complexity as Information Integration

Additionally, we are able to use this information-theoretic approach to assess the complexity of the controller. In this setting the controller is given by a multivariate random variable $C_t = (C_t^1, \dots C_t^n)$ in each point in time t, as depicted for $n = 2$ on the left hand side of Figure 8. In the previous sketches we display only one node per random vector regardless of the dimension to increase readability.

Here we deem a controller to be complex if it is more than the sum of its parts. Hence, we quantify the difference between the whole system and one in which the different parts of the controller do not interact with each other. In the latter case there is no information flow from C_t^i to C_{t+1}^j if $i \neq j$, marked here by the dashed connections in Figure 8. This means that the different controller nodes do not integrate information among each other.

On the left hand side of Figure 8 we only depict the controller without explicitly including external influences, similar to the setting in [41, 4]. Instead, we now embed this measure in the sensorimodor loop, which leads to the more complicated sketch on the right hand side of Figure 8 that depicts the following measure

$$\Phi = \sum_i I(C^i_{t+1}; C^{I\setminus i}_t | C^i_t, S_{t+1}),$$

which is discussed in the context of other information-theoretic measures in [31, 32]. Further similar complexity measures can be found in [7, 8, 24, 41, 4, 36].

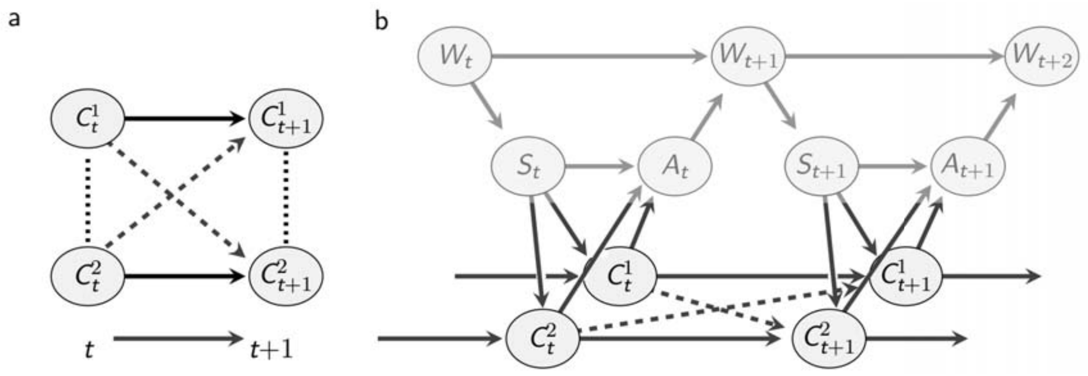

Figure 8. Sketch of the Integrated Information in the controller in (a) and the Integrated Information embedded in the sensorimotor loop in (b).

This measure of complexity is related to the Integrated Information Theory of consciousness in which the capability of a system to integrate information is equated with its amount of consciousness. This theory started as a measure for brain complexity in [55] and developed through different iteration, mainly by Tononi and his coworkers, towards a broad theory of consciousness, in, for instance, [54, 40, 2].

The different branches and iterations of the theory define various approaches to measuring the Integrated Information in a system. The information-theoretic measure discussed above can be seen in the context of [54] in which Tononi summarizes the theory as follows:

> “In short, integrated information captures the information generated by causal interactions in the whole, over and above the information generated by the parts.”

The information-theoretic framework discussed here allows us to calculate the Integrated Information of a system as well as other measures such as the Morphological Computation and relate them to each other. We discuss our previous results in the next section.

Although the Integrated Information Theory is largely restricted to only considering the controller, the following selection of studies relates the Integrated Information to an agent’s behavior and environment. In [17] the authors simulate evolving agents in a maze and observe that the Integrated Information increases with the fitness of the agents. The authors of [3] conclude that an agent needs an increased Integrated Information value in order to navigate in a complex environment and in [1] a high Integrated Information increases the likelihood of a rich, dynamical agent behavior.

7. Applications and Previous Results

Here we discuss two situations for which we have intuitive hypotheses that we were able to confirm using the information-theoretic approach. There we measure the information flows in simplistic settings with simulated agents.

When an embodied agent performs a task then it can make use of the properties of its body in a more or less optimal way, using higher or lower amounts of Morphological Computation. Hence, we would expect that the necessary controller complexity is inversely correlated with the amount of Morphological Computation, meaning the more an agent exploits its morphology, the less information integration in the brain is required.

This notion was summarized in [47], where the authors characterize the relationship between the morphology and the controller in the following way:

> "Second, there is a kind of trade-off or balance: the better the exploitation of the dynamics, the simpler the control, the less neural processing will be required."

We are able to test this hypothesis by calculating the information-theoretic measures for Morphological Computation and the controller complexity, discussed above. In the article [33] we use simple simulated agents that move inside a racetrack and have the goal to not touch the walls. These simulated agents consist of a round body and two forward facing binary sensors that can detect a wall. Additionally, they can be thought of as having two wheels that can spin either fast or slow, resulting in four different movements, as depicted on the left of Figure 9. Inside the agent are the different components, namely sensors, actuators and the controller, as described in Section 3. The controller consists in this case of two binary nodes such that the measure for the complexity of the controller corresponds to the one depicted in Figure 8. The environment of these agents, with five agents moving in it, can be seen in Figure 9 (b). In (c) there are seven different agents displayed with sensor length from 0.5 on the top to 2 on the bottom. For each sensor length we select an optimal behavior using the concept of planning as inference, discussed in, for instance, [6, 57, 56]. More details can be found in [33, 32] and the corresponding code can be accessed at [30].

Figure 10 depicts the result of the Morphological Computation and the controller complexity. The used Morphological Computation measure quantifies the information flowing through the environment, given the action, which corresponds to the connection marked by (1) in Figure 7. The x-axis of Figure 10 displays the quality of the sensors, in this case the reach of the binary sensors. Varying this has a direct influence on the interaction among an agent's body and its environment. In Figure 10 we can observe

IOP Conf. Series: Materials Science and Engineering **1321** (2024) 012008 doi:10.1088/1757-899X/1321/1/012008

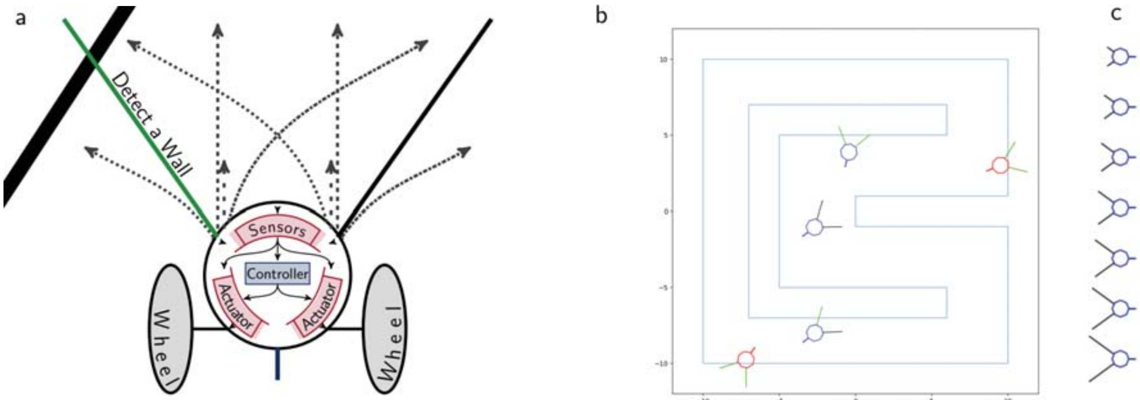

Figure 9. (a) A sketch of the two-wheeled agent with its two sensors and four different movements, (b) the environment with five different agents acting in it and (c) agents with sensor length ranging from 0.5 on the top to 2 on the bottom.

an inversely correlated relationship between the Morphological Computation and the complexity of the controller, given by the Integrated Information. The more an agent interacts with the environment, the less complexity in the controller is required, as hypothesized. So, the longer the sensors of the agents are, the more they can perceive of their environment and the better they can interact with it. This better interaction with the environment makes it possible for the agent to select appropriate actions without the need for a complex control architecture.

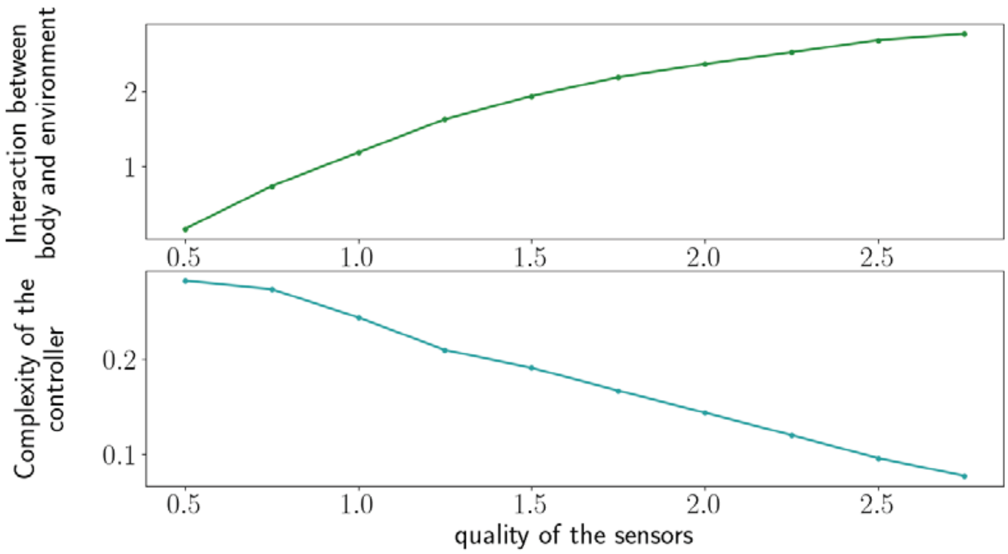

Figure 10. The arithmetic means over 100 agents of the Morphological Computation on the top and the complexity of the controller in the bottom, discussed in detail in [33].

Now the question remains why we would need a complex controller in agents that perform simple tasks at all. We could design the morphology of the agents in a way that minimizes the need for a complex controller. Naturally, we would need a higher complexity in the case of agents that are faced with more advanced settings or that have to perform multiple tasks.

Here, we want to instead highlight the challenge that learning a new task poses. As discussed in the introduction, via the example of a musician, performing a learned task is often much easier compared to acquiring a novel behavior in the first place. The agent might not know initially how to utilize the properties of the body best in order to use Morphological Computation. Hence, our hypothesis would be that a learning agent first needs a higher controller complexity, a higher integration of information, which then decreases as the model learns to exploit the properties of the body.

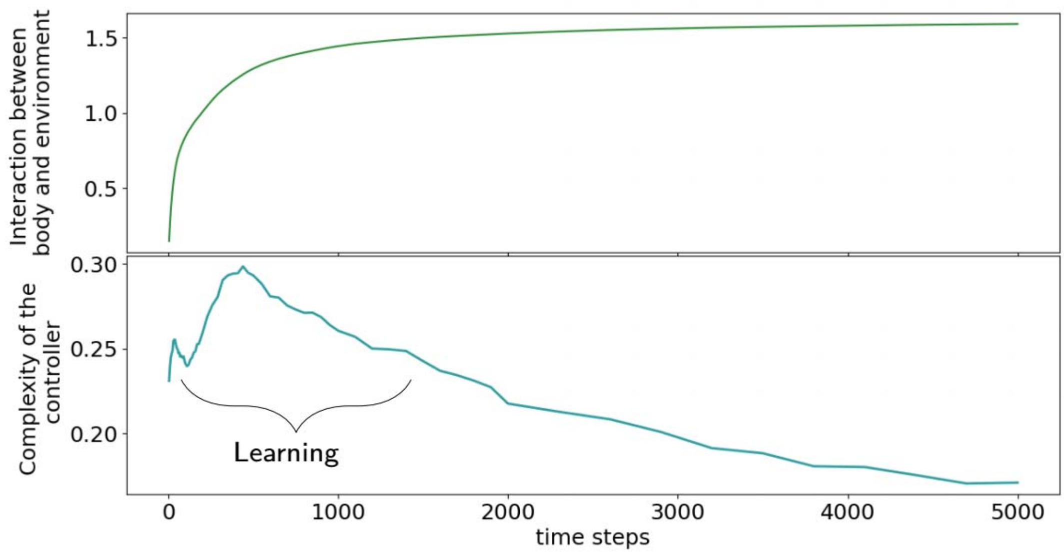

Figure 11. The results of the measures for the Morphological Computation on the top and the complexity of the controller on the bottom of learning agents, discussed in more detail in [34].

In [34] we observe simple simulated agents that learn how to move inside the racetrack depicted in Figure 9 (b) and to avoid the walls using a modified *em*-algorithm. The agents are very similar to the ones depicted in Figure 9 (a), but they additionally have an internal prediction of their next sensory state in order to enable the learning. A video of an agent in its environment and the code can be found at [29]. The learning algorithm is described in more detail in the appendix of [34] and in Section 5.4 of [32]. The results of the agents that do learn to avoid the walls are given in Figure 11. On the top we see the increasing Morphological Computation and the bottom depicts the complexity of the controller. Here the controller complexity includes an Integrated Information measure as well as a synergistic measure for the internal prediction. We observe that there is an increased controller complexity in the first approximately 1500 time steps, after which the controller complexity decreases. Note that agents which do not learn to perform the task have a lower Morphological Computation value and a constantly high controller complexity.

We conclude that an increased controller complexity is necessary for the agent to understand and exploit the environment. The agent forms an internal model of the world to predict the next sensory state. In order to learn an accurate model of the dynamics of the environment the agent needs to combine different information sources, which leads to a higher complexity of the controller. If the agent is able to form an accurate world model, then the agent can make use of its body's interaction with the environment. Hence, a higher controller complexity can facilitate a better interaction of the agent's body with its environment, leading to a higher Morphological Computation. This higher Morphological Computation then in turn reduces the necessity for a complex controller. In conclusion, Morphological Computation and the controller complexity influence each other.

8. Discussion

The information-theoretic framework discussed here enables us to measure all the information flows among the different parts of the sensorimotor loop, namely the environment and the sensor, actuator and controller nodes. Here it is possible to not only quantify the pair-wise influences, but also the synergistic impact that two or more variables have together on an additional variable.

Alternative information-theoretic approaches to measuring an information flow make use of interventional methods, as discussed in, for instance, [11, 23]. Additionally, the minimization defined here in Definition 4.1 is proposed to be changed in [9]. There the author proposes to restrict the set $\mathcal{M}$ is further to more accurately describe the underlying channel.

One caveat of the presented method is that it is only applicable for small, discrete networks at the moment. Calculating the introduced information-theoretic measures involves multiple iterations over the state spaces of the corresponding probability distributions, which is very costly for large systems. Hence, for larger networks the involved conditional probability distributions become intractable and would need to be approximated. Similarly, calculating the measures in a continuous setting is a challenge, because there the relative entropy can rarely be calculated directly but needs to be estimated. Nonetheless, extending this framework to a continuous setting would make it possible to analyze a wide range of more natural situations and is therefore a promising direction for future research.

Furthermore, the measures discussed in this setting are defined using Shannon's information theory, as introduced in Section 2. In this theory the importance of an information is solely judged by the frequency of its occurrence and not by its content, as Shannon describes in [51] as follows:

> "The fundamental problem of communication is that of reproducing at one point either exactly or approximately a message selected at another point. Frequently the messages have meaning; that is they refer to or are correlated

according to some system with certain physical or conceptual entities. These semantic aspects of communication are irrelevant to the engineering problem."

Hence, the syntax or semantics of a message are not explicitly included in the theory, but one could argue that they would implicitly emerge over time.

Nonetheless, the introduced information-theoretic measures provide us with a framework in which we can analyze every information flow among the different parts of an acting agent. As demonstrated in Section 7 this can be used to confirm our intuitions about the influences of the controller and the environment. In the future it might lead to the discovery of interesting, counter intuitive relationships among body, brain and environment.

In conclusion, information theory provides useful tools that could serve as an additional viewpoint to analyze artificial learning and acting agents.

Funding

The authors acknowledge funding by the Deutsche Forschungsgemeinschaft Priority Programme "The Active Self" (SPP 2134).

References

[1] L. Albantakis and G. Tononi. "The Intrinsic Cause-Effect Power of Discrete Dynamical Systems—From Elementary Cellular Automata to Adapting Animats". In: *Entropy* (2015), pp. 5472–5502. DOI: 10.3390/e17085472.

[2] L. Albantakis, L. Barbosa, G. Findlay, M. Grasso, A. M. Haun, W. Marshall, W. G. P. Mayner, A. Zaeemzadeh, M. Boly, B. E. Juel, S. Sasai, K. Fujii, I. David, J. Hendren, J. P. Lang, and G. Tononi. "Integrated information theory (IIT) 4.0: Formulating the properties of phenomenal existence in physical terms". In: *PLOS Computational Biology* 19.10 (Oct. 2023), pp. 1–45. DOI: 10.1371/journal.pcbi.1011465. URL: https://doi.org/10.1371/journal.pcbi.1011465.

[3] L. Albantakis, A. Hintze, C. Koch, C. Adami, and G. Tononi. "Evolution of Integrated Causal Structures in Animats Exposed to Environments of Increasing Complexity". In: *PLOS Computational Biology* 10.12 (2014), pp. 1–19. DOI: 10.1371/journal.pcbi.1003966.

[4] S. Amari, N. Tsuchiya, and M. Oizumi. "Geometry of Information Integration". In: *Information Geometry and Its Applications*. Ed. by N. Ay, P. Gibilisco, and F. Matúš. Cham: Springer, 2018, pp. 3–17. ISBN: 978-3-030-07405-0.

[5] W. R. Ashby. *An introduction to cybernetics*. New York: J. Wiley, 1956.

[6] H. Attias. "Planning by Probabilistic Inference". In: *Proceedings of the Ninth International Workshop on Artificial Intelligence and Statistics*. 2003, pp. 9–16.

[7] N. Ay. "Information Geometry on Complexity and Stochastic Interaction". In: *MPI MIS PREPRINT 95* (2001). URL: `https://www.mis.mpg.de/preprints/2001/preprint2001_95.pdf` (visited on 04/28/2022).

[8] N. Ay. "Information Geometry on Complexity and Stochastic Interaction". In: *Entropy* 17 (2015), pp. 2432–2458. DOI: `10.3390/e17042432`.

[9] N. Ay. "Confounding Ghost Channels and Causality: A New Approach to Causal Information Flows". In: *Vietnam Journal of Mathematics* (2021), pp. 547—576. DOI: `10.1007/s10013-021-00511-w`.

[10] N. Ay and W. Löhr. "The Umwelt of an embodied agent-a measure-theoretic definition". In: *Theory in Biosciences* 134 (Dec. 2015), pp. 105—116. DOI: `10.1007/s12064-015-0217-3`.

[11] N. Ay and D. Polani. "Information flows in causal networks". In: *Advances in complex systems* 11.1 (2008), pp. 17–41. ISSN: 0219-5259. DOI: `10.1142/S0219525908001465`.

[12] N. Ay and K. Zahedi. "On the Causal Structure of the Sensorimotor Loop". In: *Guided Self-Organization: Inception*. Ed. by M. Prokopenko. Berlin, Heidelberg: Springer, 2014, pp. 261–294. DOI: `10.1007/978-3-642-53734-9_9`.

[13] R. A. Brooks. "Intelligence without reason". In: *Proceedings of the 12th international joint conference on artificial intelligence (IJCAI-91)*. Ed. by R. R. Myopoulos J. Morgan Kaufmann publishers Inc., 1991, pp. 569–595.

[14] R. A. Brooks. "Intelligence without representation". In: *Artificial Intelligence* 47.1 (1991), pp. 139–159. ISSN: 0004-3702. DOI: `10.1016/0004-3702(91)90053-M`.

[15] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. New York: John Wiley & Sons, Inc., 2006. ISBN: 978-0-471-06259-2.

[16] J. Darroch and D. Ratcliff. "Generalized iterative scaling for log-linear models". In: *The Annals of Mathematical Statistics* 43 (1972), pp. 1470–1480. DOI: `10.1214/aoms/1177692379`.

[17] J. A. Edlund, N. Chaumont, A. Hintze, C. Koch, G. Tononi, and C. Adami. "Integrated Information Increases with Fitness in the Evolution of Animats". In: *PLoS Computational Biology* 7.10 (2011). DOI: `10.1371/journal.pcbi.1002236`.

[18] K. Ghazi-Zahedi. *Morphological Intelligence*. Cham: Springer, 2019. ISBN: 978-3-030-20620-8.

[19] K. Ghazi-Zahedi, D. Häufle, G. Montúfar, S. Schmitt, and N. Ay. "Evaluating Morphological Computation in Muscle and DC-Motor Driven Models of Hopping Movements". In: *Frontiers in Robotics and AI* 2 (2016). DOI: `10.3389/frobt.2016.00042`.

[20] K. Ghazi-Zahedi, C. Langer, and N. Ay. "Morphological Computation: Synergy of Body and Brain". In: *Entropy* 19.456 (2017). DOI: `10.3390/e19090456`.

[21] K. Ghazi-Zahedi and N. Ay. "Quantifying Morphological Computation". In: *Entropy* 15.5 (2013), pp. 1887–1915. DOI: `10.3390/e15051887`.

[22] K. Ghazi-Zahedi, R. Deimel, G. Montúfar, V. Wall, and O. Brock. "Morphological Computation: The Good, the Bad, and the Ugly". In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems*. 2017, pp. 464–469. DOI: `10.1109/IROS.2017.8202194`.

[23] D. Janzing, D. Balduzzi, M. Grosse-Wentrup, and B. Schölkopf. "Quantifying causal influences". In: *The Annals of Statistics* 41.5 (2013), pp. 2324 –2358. DOI: `10.1214/13-AOS1145`. URL: `https://doi.org/10.1214/13-AOS1145`.

[24] M. S. Kanwal, J. A. Grochow, and N. Ay. "Comparing Information-Theoretic Measures of Complexity in Boltzmann Machines". In: *Entropy* 19.310 (2017). DOI: `10.3390/e19070310`.

[25] D. P. Kingma and M. Welling. *Auto-Encoding Variational Bayes*. 2022. arXiv: `1312.6114 [stat.ML]`.

[26] A. S. Klyubin, D. Polani, and C. L. Nehaniv. "Representations of Space and Time in the Maximization of Information Flow in the Perception-Action Loop". In: *Neural computation* 19 (Oct. 2007), pp. 2387–432. DOI: `10.1162/neco.2007.19.9.2387`.

[27] A. S. Klyubin, D. Polani, and C. L. Nehaniv. "Tracking Information Flow through the Environment: Simple Cases of Stigmergy". In: *Artificial Life IX: Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems*. 2004, pp. 563–568. ISBN: 978-0-262-66183-6. URL: `https://uhra.herts.ac.uk/bitstream/handle/2299/3945/101985.pdf?sequence=4&isAllowed=y` (visited on 08/10/2023).

[28] A. Klyubin, D. Polani, and C. Nehaniv. "Empowerment: a universal agent-centric measure of control". In: *2005 IEEE Congress on Evolutionary Computation*. Vol. 1. 2005, pp. 128–135. DOI: `10.1109/CEC.2005.1554676`.

[29] C. Langer. *LearningRequiresIntInf GitHub Repository*. 2022. URL: `https://github.com/CarlottaLanger/LearningRequiresIntInf` (visited on 09/01/2022).

[30] C. Langer. *Morphology Shapes Integrated Information GitHub Repository*. 2021. URL: `https://github.com/CarlottaLanger/MorphologyShapesIntegratedInformation` (visited on 07/20/2023).

[31] C. Langer and N. Ay. "Complexity as Causal Information Integration". In: *Entropy* 22.10 (2020). DOI: `10.3390/e22101107`.

[32] C. Langer. "Information Integration in Embodied Systems". PhD thesis. Hamburg University of Technology, 2024.

[33] C. Langer and N. Ay. "How Morphological Computation Shapes Integrated Information in Embodied Agents". In: *Frontiers in Psychology* 12 (2021). ISSN: 1664-1078. DOI: `10.3389/fpsyg.2021.716433`.

[34] C. Langer and N. Ay. "Outsourcing Control Requires Control Complexity". In: *Artificial Life* (June 2024), pp. 1–22. ISSN: 1064-5462. DOI: `10.1162/artl_a_00443`.

[35] D. J. C. MacKay. *Information Theory, Inference and Learning Algorithms.* Cambridge: Cambridge University Press, 2003. ISBN: 978-0-521-64298-9.

[36] P. A. Mediano, A. K. Seth, and A. B. Barrett. "Measuring Integrated Information: Comparison of Candidate Measures in Theory and Simulation". In: *Entropy* 21.1 (2018). ISSN: 1099-4300. DOI: `10.3390/e21010017`.

[37] G. Montúfar, K. Ghazi-Zahedi, and N. Ay. "A Theory of Cheap Control in Embodied Systems". In: *PLOS Computational Biology* 11.9 (2015). DOI: `10.1371/journal.pcbi.1004427`.

[38] V. Müller and M. Hoffmann. "What Is Morphological Computation? On How the Body Contributes to Cognition and Control". In: *Artificial Life* 23 (2017). DOI: `10.1162/ARTL_a_00219`.

[39] K. Nakajima, H. Hauser, R. Kang, E. Guglielmino, D. Caldwell, and R. Pfeifer. "A soft body as a reservoir: case studies in a dynamic model of octopus-inspired soft robotic arm". In: *Frontiers in Computational Neuroscience* 7 (2013). DOI: `10.3389/fncom.2013.00091`.

[40] M. Oizumi, L. Albantakis, and G. Tononi. "From the Phenomenology to the Mechanisms of Consciousness: Integrated Information Theory 3.0". In: *PLOS Computational Biology* (2014). DOI: `10.1371/journal.pcbi.1003588`.

[41] M. Oizumi, N. Tsuchiya, and S. Amari. "Unified framework for information integration based on information geometry". In: *JournalProceedings of the National Academy of Sciences of the United States of America*. 2016, pp. 14817–14822. DOI: `10.1073/pnas.1603583113`.

[42] T. Parr, G. Pezzulo, and K. J. Friston. *Active inference: the free energy principle in mind, brain, and behavior*. MIT Press, •.

[43] J. Pearl. *Causality*. Cambridge: Cambridge University Press, 2009. ISBN: 978-0-511-80316-1.

[44] J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. San Francisco: Morgan Kaufmann Publishers, 1988. ISBN: 978-1-55860-479-7.

[45] J. Pearl. "On the Interpretation of do(x)". In: *Journal of Causal Inference* 7 (2019). DOI: `10.1515/jci-2019-2002`.

[46] R. Pfeifer and J. Bongard. *How the body shapes the way we think: a new view of intelligence*. Cambridge: The MIT Press (Bradford Books), 2006. ISBN: 978-0-262-53742-1.

[47] R. Pfeifer and G. Gómez. "Morphological Computation – Connecting Brain, Body, and Environment". In: *Creating Brain-Like Intelligence: From Basic Principles to Complex Intelligent Systems*. Ed. by B. Sendhoff, E. Körner, O. Sporns, H. Ritter, and K. Doya. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 66–83. DOI: `10.1007/978-3-642-00616-6_5`.

[48] R. Pfeifer, F. Iida, and G. Gomez. “Morphological computation for adaptive behavior and cognition”. In: *International Congress Series* 1291 (June 2006), pp. 22–29. DOI: 10.1016/j.ics.2005.12.080.

[49] C. Salge, C. Glackin, and D. Polani. “Empowerment–An Introduction”. In: *Guided Self-Organization: Inception*. Ed. by M. Prokopenko. Berlin, Heidelberg: Springer, 2014, pp. 67–114. DOI: 10.1007/978-3-642-53734-9_4.

[50] M. Seitzer, B. Schölkopf, and G. Martius. “Causal Influence Detection for Improving Efficiency in Reinforcement Learning”. In: *Advances in Neural Information Processing Systems*. Ed. by A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan. 2021. URL: https://openreview.net/forum?id=DXJl9826dm (visited on 07/10/2023).

[51] C. E. Shannon. “A Mathematical Theory of Communication”. In: *The Bell System Technical Journal* 27 (1948), pp. 379 –423. ISSN: 0005-8580.

[52] C. Shannon and W. Weaver. *The Mathematical Theory of Communication*. The University of Illinois Press, 1949. ISBN: 978-0-252-72548-7.

[53] N. Tishby and D. Polani. “Information Theory of Decisions and Actions”. In: *Perception-Action Cycle: Models, Architectures, and Hardware*. Ed. by V. Cutsuridis, A. Hussain, and J. G. Taylor. New York: Springer, 2011, pp. 601–636. DOI: 10.1007/978-1-4419-1452-1_19.

[54] G. Tononi. “Consciousness as Integrated Information: a Provisional Manifesto”. In: *The Biological Bulletin* 215 (2008), pp. 216–242. DOI: 10.2307/25470707.

[55] G. Tononi, O. Sporns, and G. M. Edelman. “A measure for brain complexity: Relating functional segregation and integration in the nervous system”. In: *Proceedings of the National Academy of Sciences of the United States of America* (1994), pp. 5033–5037. DOI: 10.1073/pnas.91.11.503.

[56] M. Toussaint. “Probabilistic inference as a model of planned behavior”. In: *Künstliche Intelligenz* 23.3 (2009), pp. 23–29.

[57] M. Toussaint, L. Charlin, and P. Poupart. “Hierarchical POMDP Controller Optimization by Likelihood Maximization”. In: *Proceedings of the Twenty-Fourth Conference Annual Conference on Uncertainty in Artificial Intelligence*. 2008, pp. 562–570. DOI: 10.5555/3023476.3023543.