Minimally Cognitive Robotics: Body Schema, Forward Models, and Sensorimotor Contingencies in a Quadruped Machine

Matej Hoffmann[1,2]

[1] Artificial Intelligence Laboratory, Department of Informatics, University of Zurich, Andreasstrasse 15, 8050 Zurich, Switzerland
[2] iCub Facility, Istituto Italiano di Tecnologia, Via Morego 30, 16163 Genoa, Italy,
matej.hoffmann@iit.it

Abstract. In response to the cognitivistic paradigm and its problems, the embodied cognition viewpoint was proposed. In robotics, this resulted in a radical move away from higher-level cognitive functions toward direct, almost "brain-less" interaction with the environment (e.g., behavior-based robotics). While some remarkable behaviors were demonstrated, the complexity of tasks the agents could master remained limited. A natural extension of this approach lies in letting the agents extract regularities in sensorimotor space and exploit them for more effective action guidance. We will use a collection of case studies featuring a quadruped robot to concretely explore this space of minimally cognitive phenomena and contrast the concepts of body schema, forward internal models, and sensorimotor contingencies. The results will be interpreted from a "grounded cognition" and a non-representationalist or enactive perspective. Finally, the utility of robots as cognitive science tools and their compatibility with different cognitive science paradigms will be discussed.

Keywords: developmental robotics, cognitive robotics, body schema, forward models, sensorimotor contingencies, grounded cognition, enaction, synthetic methodology, embodied cognition

1 Introduction

Within the cognitivistic paradigm in cognitive science (e.g., [19, 47]), thinking is understood as a result of computation over symbols that represent the world. On the other hand, physical activities, like walking, may be looked at as very low-level, simple and, therefore, uninteresting with regard to the study of cognition. More recently, the view of cognition as symbolic computation has been challenged, and an embodied, action-oriented, dynamic, and developmental view has been proposed instead (e.g., [58, 55, 39, 45, 16]). The boundaries between cognitive and non-cognitive phenomena have started to blur and the key influence of the body and the physical interaction with the environment has become accepted. Furthermore, a central role of developmental processes in the emergence of cognition has been asserted. There is growing and increasingly detailed evidence from psychology and neurosciences in support of the *embodied cognitive science* view. However, the premises of the new paradigm—whole brain-body-environment systems rather than isolated subsystems should be studied over extended time periods—pose new challenges to practical empirical research in animals and humans. Here, cognitive developmental robotics as a synthetic approach, i.e. instantiating and studying the phenomena of interest in robots, can serve as a viable tool to verify certain hypotheses and complement the research in psychology and neuroscience [44, 3, 42].

In this article[3], we will first try to categorize the research in robotics from an (embodied) cognitive science viewpoint (Section 2). Then we will use a quadruped robot and investigate the possibilities of its autonomous development from simple reactive to the first cognitive behaviors: from locomotion to cognition (Section 3). The scenarios are chosen such that they can be successfully mastered only if the robot leaves the "here-and-now" time scale of reactive, stimulus-response, behaviors [59]. In order to do that, the robot needs to extract some regularities from its interaction with the environment and utilize them when selecting the next actions to take. In Section 4, we will then attempt to interpret the case studies from two different viewpoints: a grounded cognition [4] or minimal robust representationalist [9] perspective followed by a non-representationalist or enactive account [58, 54]. In the case studies, we explored three concepts that were proposed to explain the development and operation of minimal instances of cognition: body schema (e.g., [29, 12]), forward internal models (e.g., [61, 11]), and sensorimotor contingencies (SMCs) [39]. Concrete implementations in the robot help us to better understand the meaning of each of them and implications for cognitive development—this analysis will be the topic of Section 5. Then, the implications and limitations of using robots as cognitive science tools will be discussed (Section 6). In particular, we will investigate, whether robots can serve as an essentially paradigm-neutral research tool, or whether their use poses intrinsic limitations—with regard to enactive cognitive science viewpoint, for example. We will close with a conclusion.

2 A Cognitive Classification of AI and Robotics Research

In this section we will strive to sketch a "cognitive landscape" in order to classify some seminal work in Artificial Intelligence (AI) and robotics from the point of view of cognitive science. Obviously, cognition is a very difficult phenomenon and any attempt to "pin it down" in a single diagram is bound to fail. Nevertheless, we believe that it will still be valuable to depict some of the key facets of cognition and their instantiation in AI and robotics in graphical form. To this end, we propose four different axes and hence two different 2D schematics.

[3] Parts of this article are based on the author's PhD thesis [28].

2.1 Offline Reasoning Capability vs. Real-Time Responsiveness

In the first diagram (Fig. 1), the y-axis essentially follows the "grounded cognition" viewpoint of Barsalou [4], and Clark & Grush [9], who used the capability of offline reasoning (or "environmentally decoupled thought") to demarcate cognitive agents (from non-cognitive agents). However, if the role of cognition is to support purposeful and timely action in the real world, the "cognitive" axis needs to be complemented by another dimension, which we have labeled "real-time responsiveness". In other words, cognition should not come at the expense of fast interaction with the environment.

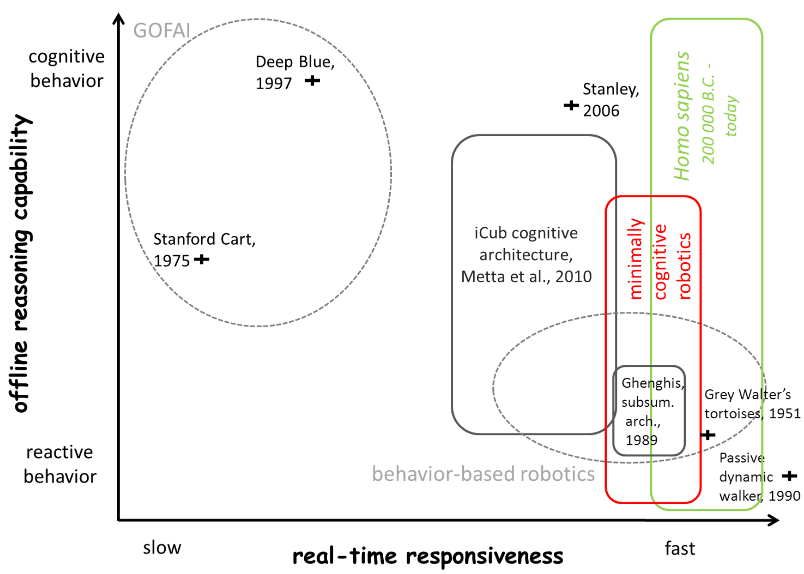

Fig. 1. Cognitive landscape - Offline reasoning capability vs. real-time responsiveness. This figure attempts to classify selected work in AI and robotics from the perspective of cognitive science according to the agents' capability of "running cognition offline" vs. preserving the means to respond immediately. See text for details.

In most machines, their creators decided for a single intersection of these two dimensions—that is the control architecture operates on only one level. These are represented by crosses on the schematics. Passive dynamic walkers [36] can be depicted at the very bottom of the "cognitive axis" (y-axis) as they are passive mechanical machines and are completely coupled to their physical environment. At the same time, their real-time responsiveness is maximum. Examples of reactive agents—creatures capable of simple stimulus-response behavior only—would also occupy the bottom of the "cognitive axis". The tortoises of Grey Walter [60] were composed of direct, analog links connecting sensors and motors; consequently, the real-time responsiveness of the tortoises will be almost as high as that of the passive dynamic walkers. Examples of Good-Old Fashioned Artificial Intelligence (GOFAI), on the other hand, end up on the other end of spectrum: The chess computer, Deep Blue, is definitely capable of offline reasoning (pondering thousands of hypothetical evolutions of the game) and it is subject to soft time constraints—it does not have to respond immediately, yet its total thinking time is limited. From mobile robots, the Stanford Cart [38] was capable of offline reasoning, yet to the extent that it had lost real-time responsiveness (thinking around 15 minutes before every 1 m lurch)[4]. A modern sibling of the Cart, Stanley (DARPA Grand Challenge winner, [57]), is capable of autonomously planning and following a path in an outdoor environment, while preserving very good responsiveness to the current situation on the road.

However, a single degree of "offline reasoning capability"—cognitive, reactive, or only physical—is not sufficient to master a variety of capacities that more complex organisms demonstrate. Therefore, they typically employ a weak hierarchy of different levels ranging from mechanical feedback loops (e.g., [6]) over simple spinal reflexes, which involve direct sensorimotor connections, to more complex and abstracted layers that are present in the brain.

To illustrate this, we have depicted humans ("Homo sapiens") with a large region which ranges all the way from reactive to cognitive behavior on the "cognitive axis". Another example of behavior-based robotics [2], next to Walter's tortoises, is the robot Ghenghis [7]. The so-called subsumption architecture consists of different layers—all of them essentially reactive. Hence, it is also depicted with a small region in the reactive domain rather than a single cross. The cognitive architecture of the iCub humanoid robot as presented in [37] also contains "reactive layers", but at the same time, certain modules reach out of the simple stimulus-response realm to more decoupled processing. Both—Stanley and iCub—are capable of offline reasoning, while preserving real-time responsiveness. Alas, considerable computational resources are required. Finally, "minimally cognitive robotics"—the focus of our interest—would correspond to less "offline reasoning" and more "real-time responsiveness" than Stanley or iCub, building directly on top of the behavior-based robotics school.

2.2 Nature of "Neural Vehicles" and Their Plasticity

In the second diagram (Fig. 2), we propose two additional axes. The y-axis depicts the nature of the internal informational structures that mediate the agent's interaction with the world. They were called "neural vehicles" by Engel [15], avoiding the problematic label of "representation" ("neural" is not be taken literally and is synonymous with internal or belonging to the controller; more details will be provided in Section 4.2). The axis spans the space from no neural vehicles over internal structures operating on the sensorimotor space to symbolic spaces. The x-axis characterizes the degree to which the system has

[4] It should be noted that this was largely due to the computational power available at that time.

been engineered and remains fixed afterwards or—at the left end of the axis—how much was learned without prior knowledge and how adaptive or plastic the neural vehicles are.

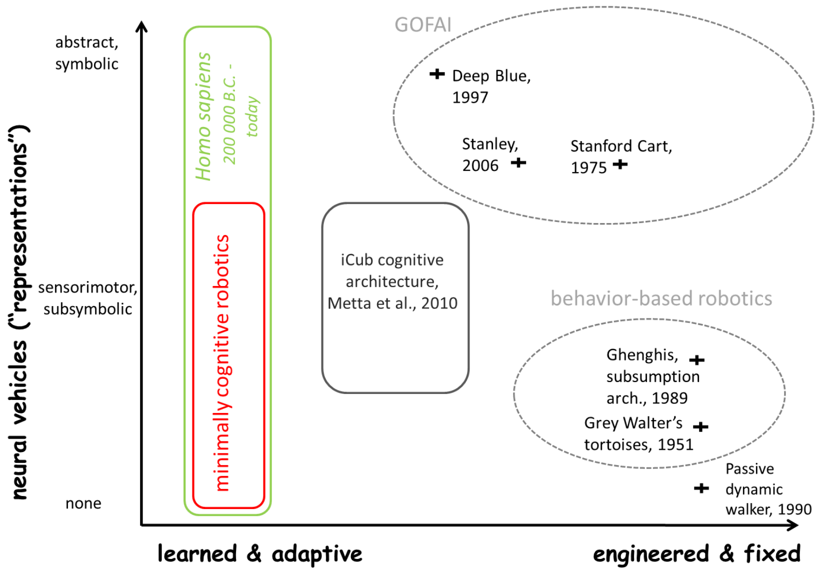

Fig. 2. Cognitive landscape - nature of "representations" vs. their plasticity. This figure attempts to classify selected work in AI and robotics using two additional axes compared to Fig. 1: the character of the control structures vs. their adaptivity. See text for details.

Humans are presumably relying on a multitude of mechanisms from direct physical interaction with the world, to simple reflexes and spinal control, to neural vehicles that operate on highly abstracted levels. Hence, they essentially span the whole range of the y-axis. Their location on the x-axis depends on the stance one takes in the "nature vs. nurture" debate. Some capabilities were learned on the evolutionary time scale, which—from the point of view of the individual—represent the "design"/engineering. Yet, the adaptive capacities of humans are extraordinary and support the positioning toward the left end of the axis.

Obviously, there are no neural vehicles, or representations, in the passive dynamic walker. The other instances of behavior-based robotics rely on simple connections of sensors and motors and are thus approaching the "sensorimotor region" from the bottom. On the x-axis, they all have the same value, as they were simply designed by their creators. The GOFAI examples fall on the opposite end of the "representation" axis, since they rely on very abstract or symbolic internal structures. On the "learned vs. engineered" axis, the mobile robots Cart and Stanley display only a limited degree of plasticity (data-driven parameter tuning was applied in [57]). Deep Blue is the most adaptive in this respect, as it learned the position evaluation function itself from thousands of grandmaster games. The iCub's cognitive architecture [37] consists of numerous modules, the core competences being centered around the sensorimotor level. The modules feature different degrees of engineering vs. learning. Finally, the space that is uninhabited by robots so far and that lies at the center of our interest in this article is the "minimally cognitive robotics" region. We will present case studies that attempt to "colonize" this space by learning basic sensorimotor capacities, including integration of information over time and its deployment, starting from minimal prior knowledge.

2.3 Cognitive Developmental Robotics

The iCub humanoid robot and its positioning in our diagrams is a representative example of cognitive developmental robotics, which can be, for example, defined as follows:

> Cognitive developmental robotics (CDR) aims to provide new understanding of how human higher cognitive functions develop by means of a synthetic approach that developmentally constructs cognitive functions. The core idea of CDR is "physical embodiment" that enables information structuring through interactions with the environment, including other agents. The idea is based on the hypothesized developmental model of human cognitive functions from body representation to social behavior. [3]

CDR is thus a subset of developmental robotics in general, which has the same mission, but is not concerned with cognitive phenomena only (however, we have to keep in mind that the boundary between sensorimotor and cognitive phenomena is blurred). A review of developmental robotics is provided by Lungarella et al. [32] or by a special issue of the Infant and Child Development Journal [46]. A review of CDR is provided by Asada et al. [3].

What we have labeled "minimally cognitive robotics" can be seen as a special subset of developmental and cognitive developmental robotics that is specifically concerned with minimal settings where the first instances of offline reasoning—or learning from experience to avoid commitment to the representationalist standpoint—capabilities emerge. In the next section, we provide an overview of experiments that we have performed in a quadruped robot that demonstrate such a developmental pathway.

3 A Developmental Pathway in a Quadruped Robot

In this section, we will present a selection of the results we obtained by instantiating a "cognitive development pathway" in a quadruped robot. The case studies presented feature the key ingredients that are believed to be necessary for cognition to emerge: rich body dynamics and physical interaction with different environments, active generation of multimodal sensory stimulation and learning from this experience over different time scales. In a first study (not reported here, [26]), the robot first learned coordinated movement commands (gaits somewhat resembling those seen in nature – walk, bound, pace etc.) which later formed its motor repertoire. The other case studies deal with the sensorimotor space and the possibility for the robot to extract regularities in it and later exploit this experience in accordance with its goals. More details can be found in [28] and in individual articles reporting the results.

3.1 Specifics of Cognition in a Quadruped Robot

The main platform in our work was the quadruped robot Puppy (Fig. 3). An obvious implication of the embodied cognition stance is that the kind of cognition that will emerge will be highly dependent on the body of the agent, its sensorimotor apparatus and the environment it is interacting with. In our case, the multimodal sensory set together with the nonlinear, partly passive, dynamics of the body can be exploited to extract information about the body itself and the environment. In addition, the absence of distal sensors (camera) forces the robot to use all the modalities by actively probing the environment, which is in accordance with the action-based view on perception and cognition.

The locomotion context is particularly suited for understanding minimally cognitive behavior. Whereas "manual cognition", i.e. reaching, grasping and dexterous manipulation, is largely restricted to humans and primates, "locomotor cognition", on the other hand, can be found in much lower animals. For example, path integration was discovered in ants [64]; prediction was demonstrated in motor preparation of prey-catching behavior of a jumping spider [52]; frogs were found to be able to predict whether an aperture could be passed [10]; finally, rats were found covertly comparing alternative paths in a T-maze, thus "planning in simulation" [23] (see [41] for a review). In this work, we will present the robot with similar scenarios: path integration, terrain discrimination and gait selection, and catching another robot.

3.2 Extracting a Body Schema from Raw Sensorimotor Data

In the first study (for details see [51]), we let the robot apply different motor patterns and recorded the corresponding sensory stimulations from its multimodal sensory set comprising primarily tactile and proprioceptive sensors. Then, we systematically analyzed the directed information flows between motors and sensors and showed how the robot could infer a primitive map of its body by extracting the structure of the sensorimotor space that is invariant to changes of the controller: A random set of motor commands proved the most effective in this respect. An information theoretic method that quantifies directed information flows between two variables (sensory and motor time series in our case), transfer entropy, was used. The result is depicted in Fig. 4.

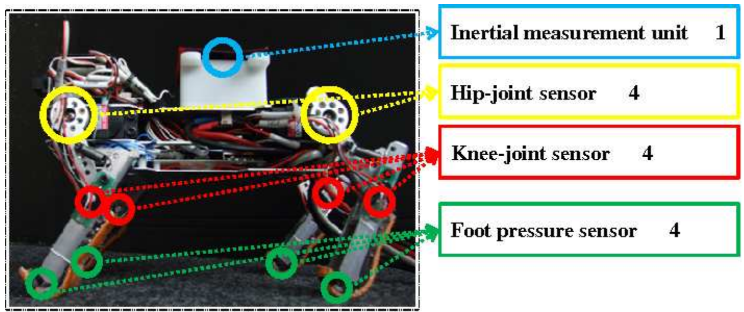

Fig. 3. The quadruped robot Puppy and its sensors. It has four active revolute joints controlled by servomotors (”shoulders“ and ”hips“, in what follows simply hips) and four passive revolute joints at the ”elbows“ and ”knees“ (later simply knees)—the passive joints have springs attached across them, making them compliant. There are angular position sensors in all the joints. In addition, there are pressure sensors on the robots feet and an inertial measurement unit (IMU – 3-axis accelerometer, 3-axis gyroscope) on the back. The robot is 20 cm long. **Labeling of channels** (to be used below). The four legs are abbreviated as FL: fore left, FR: fore right, HL: hind left, and HR: hind right. Then, M_{FL}, M_{FR}, M_{HL}, M_{HR} correspond to the four motor channels; H_{FL}, H_{FR}, H_{HL}, H_{HR} denote potentiometers in the hip joints, and K_{FL}, K_{FR}, K_{HL}, K_{HR} in the passive knee joints; P_{FL}, P_{FR}, P_{HL}, P_{HR} are feet pressure sensors, A_X, A_Y, A_Z linear accelerations in three axes, and G_X, G_Y, G_Z are angular velocities. (Figure adapted from [51].)

Unlike the majority of work on automatic model acquisition in robotics—reviewed in [24]—, which typically builds on significant prior knowledge and only refines an existing representation using vision, our method is purely data-driven and extracts the regularities intrinsic to the robot's morphology from scratch. Furthermore, the same approach can be used to move from an initial “synesthetic state”—with undifferentiated sensory modalities—to an unsupervised discovery that there are qualitatively different types of sensors.

– **Proprioceptive vs. exteroceptive modality as a graded distinction.** First, we looked specifically at information flows from motor to sensory channels. Those channels that receive strong directed information from the motor signals can be said to be “controllable” by the robot and thus reflecting the state of the the body (under the interpretation “my body is what is under my control”). Hence, these sensors can be said to have “proprioceptive” properties. Exteroceptors, on the other hand, can be defined as sensory channels sensitive to environmental changes.[5] Applying these definitions to the information flows that the agent measured gives a graded distinction of the sensors (see Fig. 5 (left)). Interestingly, only the angular position sensors in the motor-driven hip joints fell clearly into the “proprioceptive” region.

[5] We have systematically varied the environmental conditions—grounds of different friction—and analyzed the data. Please see [51] for the details.

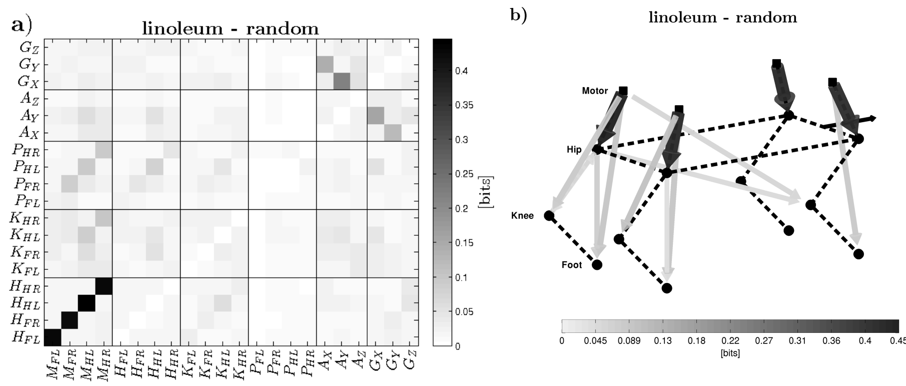

Fig. 4. Transfer entropy TE between all pairs of motor and sensory channels using random motor commands on linoleum ground. (a) Every cell of the matrix corresponds to the information transfer from the signal on the column position to the signal on the row position. Cf. Fig. 3 for the labeling of channels. **(b)** A schematic of the Puppy robot (dashed lines) with overlaid arrows depicting the TE between the individual components. For readability, only the 15 highest values are shown and the accelerometers and gyroscopes were excluded from this visualization. The strength of the information transfer is encoded as thickness of the arrows. The strongest information transfer occurs from the motor signals to their respective hip joint angles ($M_{FL} \rightarrow H_{FL}$, $M_{FR} \rightarrow H_{FR}$, $M_{HL} \rightarrow H_{HL}$, $M_{HR} \rightarrow H_{HR}$). The motors directly drive the respective hip joints and, despite some delay and noise, the hip joints always follow the motor commands, which induces a strong informational relation. The motors further show a smaller influence on the knee angles (especially at the hind legs K_{HL} and K_{HR}) and on the feet pressure sensors, all on the respective leg where the motor is mounted, thus illustrating that body topology was successfully extracted (Figure from [51].)

The other sensors—most of which would be labeled as proprioceptors using a standard "textbook" definition—were found to be more sensitive to the environment.

– **Learning about different sensory modalities.** According to O'Regan and Noe [39], it is the SMCs, i.e. the structure of the rules governing the sensory changes produced by various motor actions, what differentiates modalities. We have applied a similarity measure to the information flows and projected the sensors and motors to a 2D space, creating a sensoritopic map. The resulting map (Fig. 5 (right)) shows a reasonable clustering of angular sensors in active vs. passive joints, pressure sensors, and inertial sensors—reconfirming the SMC hypothesis and demonstrating that no additional knowledge is necessary. The motor modality, which has a different "causal content", is completely separated out on the right of the map.

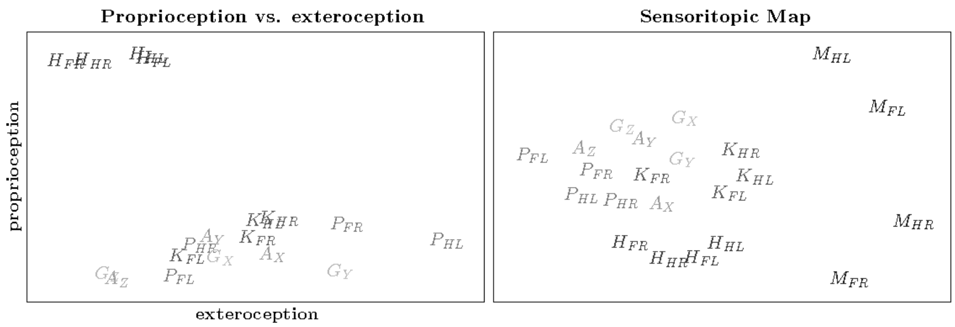

Fig. 5. Sensor spaces in the Puppy robot. (left) Proprioception vs. exteroception. **(right)** A Sensoritopic map. Projection of the sensors and motors into 2D space using multidimensional scaling based on an information flow-based similarity measure. Cf. Fig. 3 for the labeling of channels. (Figure from [51].)

3.3 Learning from Sensorimotor Experience

Whereas the previous case study had an analytical focus—how salient relationships in the sensorimotor space can be extracted—, the next logical step on the "cognitive ladder“ is to take the agent's perspective and study how it can integrate its experience and use it to improve behavior. We conducted three studies in this direction, presenting the robot with different tasks that can be successfully mastered only if the robot learns from past sensorimotor experience.

Path integration using self-motion cues. Humans, other mammals, and also arthropods are reported to be able to perform path integration: estimating the distance traveled without relying on an external reference [17, 14, 64, 65]. Odometers (step integrators) were found to play an important part in this capability. To estimate the length of the step (or stride), the animal seems to require a body representation of some sort ([65] mention: knowledge about intrinsic dynamics of limb segment motion, relationships between gait parameters and body proportions).

In our quadruped robot, we developed one possible solution to the problem: an implicit (data-driven, black-box) model that linearly combines features from multiple sensors from the robot's legs to a stride length estimate [48, 49]. Sensory features that correlated most strongly with stride length were selected and a linear regression function that combined them into a stride length estimate was derived, giving rise to a multimodal legged odometer. That is, we showed an example of a procedure that can be employed by an autonomous agent: investigate relationships between a variable of interest and the sensory (or sensorimotor) space, select the signals with the strongest relationships, and work them out into a function. The stride length estimates can then be aggregated over time, giving rise to a measure of distance traveled by the agent—a first example of *integration of information over time* in our agent.

Using sensorimotor contingencies for terrain discrimination and adaptive walking. In this study [27], a record of *past experience in the sensorimotor space was used to inform action selection*: the robot learned to estimate the effects of the application of different gaits in different contexts and used this information to choose the actions that maximize a reward signal (fast and stable locomotion). Eventually, it learned to select an appropriate subset of gaits in different contexts (see [27] for details). No abstraction or hierarchy was used, but a memory of almost raw sensorimotor sequences (compressed into features) allowed the robot to detect familiar contexts and select actions accordingly. Furthermore, we want to highlight two additional outcomes of this study:

– **Perceptual categorization from sensorimotor sequences.** Perceptual categorization can be simplified through embodied interaction with the environment and active generation of sensory stimuli (see e.g., [43]). In our study, when the robot was running on different grounds, only certain, prestructured, stimuli were inevitably induced in the sensory modalities. In addition, the particular action used at every moment—the gait—co-determined what was sensed. We demonstrate this effect by showing the improvement in ground classification when data generated by different gaits are classified separately. Furthermore, we again confirm the hypothesis (put forth in [39] and tested in a simple robot in [35]) that object categorization (the ground being the object here) is improved if longer sensorimotor sequences are considered. The data from both real and simulated robot convincingly demonstrate this (Fig. 6).

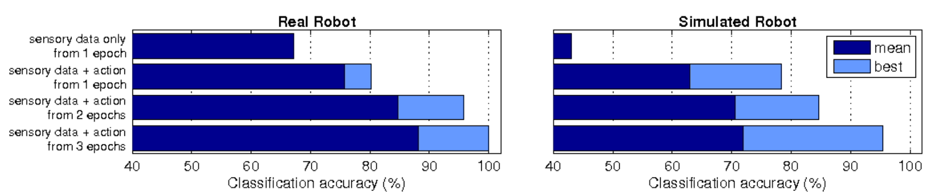

Fig. 6. Comparison of ground classification accuracies when the action context is taken into account to different degrees. The first row corresponds to data from one sensory epoch collapsed across all gaits, i.e. without the action context. Subsequent rows report results where classification was performed separately for each gait and increasingly longer histories were available. "Mean" values represent the mean performance over the individual classifications runs preconditioned on the gait the robot was using; "best" are classification results from the gait that facilitated perception the most. (left) Real robot. (right) Simulated robot. (Figure and caption from [27].)

– **Compression of sensorimotor space through embodiment.** We employed a model presented in [35] and adapted it to our situation. An exhaustive approach to remembering sensorimotor experience was used: the agent did not try to explicitly extract the structure of the sensorimotor space and store it in a compressed form; instead, every new action-observation combination and their history of up to 4 epochs (10 seconds in total) was added to the memory. Although the theoretical dimension of the sensorimotor space was enormous, due to the constraints imposed by the morphology of the robot's mechanical and sensory system, the nature of the interaction with the environment, the action repertoire, and the action selection algorithm, only a small portion of the theoretical state space was visited (2 to 4% of possible states; see [28] for the details). This is in accordance with previous findings on how sensorimotor information is structured through embodiment [33]. That is, the regularities in the sensorimotor space assist the robot in dealing with the curse of dimensionality.

Moving target seeking with forward and inverse models. This study [40] constitutes our last step of incremental cognitive development in a mobile robot. We prepared a scenario in which a "hunter" robot needs to catch its conspecific "prey" robot[6]. The scenario was manipulated in order to investigate under which conditions more elaborate planning becomes necessary and what are the best candidates for the implementation. The "hunter" robot was progressively forced by the task-environment to employ less reactive and more cognitive strategies. Finally, it arrived at a multi-step planning architecture: *a "decoupled" forward model, which can be executed independently*. This corresponds to the "cognitive hallmark" proposed by Clark and Grush [9]. The specific points addressed were:

– **Learning a forward model.** A forward model predicting the robot's change in position and orientation was learned through random exploration of the effects of different gaits. An egocentric reference frame was used and no prior knowledge about the platform (such as its kinematics or dynamics) were necessary.
– **Goal state and inverse modeling.** In order to reach the goal state—coming as close as possible to the prey robot—, an inverse model became necessary. That is, given a current position and orientation and a desired one, the output was the best action to take. This was obtained through simple Bayesian inference.
– **Multi-step planning.** We presented the robot with different scenarios: whereas a simple application of the inverse model yielded satisfactory results in some scenarios (hunter and prey in a wall-enclosed arena), in others (open environment) it did not suffice. There, we studied how multistep planning

[6] In this study, simulated Khepera robots with a discrete "gait repertoire" to mimic the situation in the quadruped robot were used. Perception of the hunter's and prey's position were simplified and "GPS" signal available in the simulator was used. Experiments on simulated models of the Puppy robot are under way.

can improve the results. In order to cope with a combinatorial explosion of possibilities, a heuristic best-first-search was implemented.

– **Extending modeling to other agents.** Finally, the utility of explicitly modeling a part of the environment (the "prey" agent) was evaluated and successfully incorporated when it improved the agent's performance. In this way, the agent extended the space of its "cognitive processes" to other agents.

4 Two Sides of the Same Coin? A "Grounded Representation" vs. a Non-representationalist Perspective on the Case Studies

In a nutshell, the case studies presented are concerned with the structure of the sensorimotor space: How it is shaped by an agent's body and dynamic interaction with the environment and how invariant relationships can be extracted and exploited by the agent to improve its behavior. While, on one hand, the performance of the robots in the tasks "speaks for itself", there are still many conceptual questions pending. In particular, should extracting and exploiting past sensorimotor experience be equated with the notions of storage, knowledge, representation, or offline reasoning? In this section, we will attempt something that is rarely undertaken: We will interpret the very same results from two different perspectives: one that posits representations followed by one that rejects them.

4.1 Increasing the Offline Reasoning Capability from Bottom-up: a Minimally Representationalist Account

The case studies presented lend themselves easily to an interpretation along the lines of "grounded cognition" [4] and "minimal robust representationalism" that was proposed by Clark & Grush [9] in defense of the notion of representation in cognitive science as well as robotics. This view essentially suggests that through processes of internalization and decoupling, cognition can eventually "run in the brain" [9, 63].

The "Extracting a body schema" case study (Section 3.2) can be naturally viewed as the robot building a sensorimotor representation of its body—a correspondence between the structure it learns (the representation) and the physical and sensorimotor properties of the robot (what is being represented) can be established. This primitive, sensorimotor structure is extracted solely from the sensory and motor channels and is thus automatically grounded. In the "Path integration" study, the robot builds what could be called a "locomotor body schema", i.e. a model of how much distance it covers every stride. Yet, this "legged odometer" was tuned by using an external reference frame; thus, the "grounding" of this representation—position and orientation in a Cartesian reference frame—is mainly on the side of the observer. The "Terrain discrimination and adaptive walking study" lends itself to a representational interpretation too. The sensorimotor histories that are stored in the associative memory ("the brain" of the agent) can be looked at as knowledge or representation of the robot's previous interactions with the world. They can also be used to classify the environments, replayed or iterated forward to get predictions and inform action selection—in accordance with the "offline cognition" notion. Finally, the "Moving target seeking" study serves as a perfect example of a bottom-up development of an internal simulator/emulator [9, 22]: the robot learns a forward and inverse model of the outcome of applying different motor commands. Later—in order to succeed in the task of catching its conspecific—it also learns a model of the "prey's" behavior and applies a multi-step planning algorithm. This demonstrates an increasing degree of offline reasoning and matches with the evolutionarily plausible path how internal representation (in the form of emulator circuitry) could possibly get its foot in the door of real-world, real-time cognition [9].

4.2 Enactivism and "Cognition-is-not-in-the-Brain" Viewpoint

Interestingly, we can try to embrace the very same case studies into a more radical school of thought that rejects the neurocentric perspective on cognition altogether. A unique perspective on cognition has been offered by the community that has grown around the work of Francisco Varela (e.g., [58]). The proponents of the *enactive framework* reject the idea that "cognition often proceeds independently of the body"[4]. For the "enactivists", cognition is not only shaped by the body and its action possibilities, but *cognition is action*—embodied action, a form of practice itself [58]. In this view, cognition is not about world-mirroring through representations, but "world-making" and sense-making. The interested reader is referred to the abundant literature (e.g., the recent collection of papers in [54], reviewed in [20]).

We will borrow useful terms from Engel et al. [15, 16] who provide a review of a turn toward action in cognitive science and propose the term *dynamic directive* to "denote the action-related role of large-scale dynamic interaction patterns that emerge in a cognitive system. On this account, directives can be defined as dispositions for action embodied in dynamic activity patterns." Importantly, the directives are not equal to states in the brain (and thus are not equal to action-oriented representations; see also [30] for a detailed philosophical account of this distinction), but refer to dynamics of the whole—or relevant parts of a—brain-body-environment system. At the same time, it may be convenient to invoke a term for the "traces" of the directives in the brain: These are the *neural vehicles* of the directives [15].

Some of the results we have presented are not compatible with this viewpoint. For example, the path integration case study was devoted to the learning of a position and orientation estimation module trained by an external reference. Moreover, in this particular study, no action selection was performed based on the path integration results. Thus, this task had little significance or meaning for the robot. In summary, the focus was on a veridical representation of the position of the robot in the environment—an emphasis that is incompatible with the formulations that belong to enactive cognitive science.

Let us look at the "adaptive walking" and the "moving target seeking" studies. There, the robot had to optimize its behavior on a task—fast and stable walking in the former case, catching another robot in the latter. To this end, different control architectures that could assist the robot in the task were explored. In the predator-prey scenario, the "neural vehicles" were data-driven and learned *ab initio*, but the structure of the model (a simple Bayesian network), the variables of interest (distances and angles), and the goal (catching the prey) came from the designers. The "world-making" of the robot has thus been relatively strongly constrained and imposed on the robot from the outside. Finally, the "adaptive walking" study, where a model of sensorimotor contingencies is employed, could probably be most in line with an enactive viewpoint. The robot simply records past sensorimotor experiences (the gait used and all the sensory channels compressed into features) together with the values of the reward function and uses this information to inform future decision-making: selecting the gait that is most likely to succeed on a given ground. The "neural vehicle" thus contains raw sensorimotor "footprints" of the robot's interaction with the environments and uses them for action guidance. The individual terrains are nowhere explicitly coded in the neural substrate—they are implicitly recognized by selecting appropriate actions. Yet, the reward function was again defined from the outside and the "sensorimotor look-up table" that is driving the behavior at discrete time steps is perhaps still too decoupled from the dynamics of the body and environment when compared to the—alas much simpler—dynamical accounts of active categorical perception [5, 8][7].

5 Body Schema, Forward Models, and Sensorimotor Contingencies: On their Overlap, Definition, and Degree of Representational Nature

We have set out to investigate bottom-up development of minimally cognitive abilities. On this path, we have repeatedly encountered three concepts: body schema, forward internal models, and sensorimotor contingencies (SMCs). We have explored them in different disguises in our robotic case studies. What can we now say regarding their nature, utility and compatibility with different cognitive science paradigms?

As reported by Rochat [50], infants spend substantial time in their early months observing and touching themselves. Through this process of babbling, intermodal redundancies, temporal contingencies and spatial congruences are picked up. This basically encompasses all the low-level relationships that an agent can learn during its early development. However, this space is too large. Therefore, in order to bootstrap its development, an agent needs to focus on some subspaces of the sensory-motor-time space and choose an appropriate way

[7] Note that the perceptual categorization we performed in Section 3.3 used solely the sensorimotor memory, i.e. the neural structure. Beer or Buhrmann et al. [5, 8], on the other hand, show that in their examples, this is not possible.

of modifying its internal dynamics in accordance with these regularities and its goals. The three aforementioned concepts qualify as suitable candidates in this regard.

5.1 The "Minimally Cognitive Concepts" in the Case Studies

- **Body schema.** As we have argued, the body has a key influence on the agent's behavior as well as on the information that enters its brain/controller (see [25] for a collection of examples illustrating this). Therefore, it can bring advantage to the agent if it can pick up the regularities that are induced by its body. The concepts of body schema and body image are used in this context. However, at the moment, they serve more as "umbrella concepts" for a multitude of body representations that animals and humans develop and use (cf. e.g. [12]). The synthetic approach allowed us to explore these concepts in more concrete terms. In the "Extracting a body schema" study (Section 3.2), we investigated two possibilities for the formation of a primitive body representation in a robot. First, we studied the structure of the sensorimotor space that is invariant to changes in the motor commands and the environment—that is, body as the invariant structure in sensorimotor space. Second, we studied which sensory channels were strongly affected by the motor signals. This provides an alternative view: the agent's body is what it can control. Both viewpoints can have merits for the agent: the former one could be used for self-diagnosis (if the invariant structure changes, this can be attributed to changes in the body), the latter one can be used to bootstrap development—learning the first behaviors. Yet, this is just the very beginning and subsequent development needs to be demonstrated. A more narrow type of body schema or image devoted specifically to estimating the robot's stride length was developed in the "Path integration" case study.
- **Forward model.** Forward model is another type of mapping that can be useful to the agent. It can be used to predict the next sensory state (given the current state and a motor command) or—if chained or iterated—even to simulate whole sensorimotor loops covertly. It is concretely defined[8] and can be instantiated at any abstraction level (i.e., not only for low-level motor control, where the existence of forward internal models is subject to a heated debate – cf. for example [11] vs. [18]). We have explicitly employed probabilistic forward and inverse models in the "Moving target seeking" study. The architecture used in the "terrain discrimination and adaptive walking" study that is using conditional probability distributions [35] also encompasses forward model functionality.
- **Sensorimotor contingencies.** Sensorimotor contingencies (SMCs) were originally presented in the influential article by O'Regan & Noe [39] as the

[8] The forward model is classically thought of as a function, $f(s_t, m) = s_{t+1}$, which maps a sensory state and a motor command to a next sensory state (where the states can be multidimensional).

structure of the rules governing sensory changes produced by various motor actions. Similarly to a body schema, this notion is still not articulated concretely enough to allow for an implementation in a robot. For example, is a forward model an instance of an SMC? Also, what is the "site" where the SMCs reside—are they stored in the brain? Very recently, Buhrmann et al. [8] have addressed these questions and proposed a dynamical systems account of SMCs, distinguishing between *Sensorimotor (SM) environment*, *SM habitat*, *SM coordination*, and *SM strategy*. The SM environment is the relation between motor actions and changes in sensory states, independently of the agent's internal (neural) dynamics. Interestingly, this definition closely resembles the forward model that we have encountered before.[9] The other notions—from SM habitat to SM strategies—add internal ("brain") dynamics to the picture. SM habitat refers to trajectories in the sensorimotor state space, but under certain conditions on the internal dynamics that is responsible for generating the motor commands. These are thus not random anymore and may depend on previous sensory states as well—an example of closed-loop control. SM coordination then further reduces the set of possible SM trajectories to those "that occur reliably and contribute functionally to a task". For example, specific patterns of pressing an object in order to assess its hardness would be SM coordination patterns serving object discrimination. Finally, SM strategies take, in addition, a normative framework ("reward" or "value" for the agent) into account.

Taking advantage of this operationalization of the SMC concept, in what disguises can we find SMCs in our case studies? In the study described in Section 3.2[10] random motor commands were applied (hence there was random or no neural dynamics) and the relationships between motor and sensory variables were studied, closely resembling the notion of SM environment.[11] Then, we also studied the relationships in the sensorimotor space when the robot was running with certain coordinated movement patterns: gaits. These were obtained by optimizing the robot's performance for speed

[9] The functional form was provided in the previous footnote. However, if the sensory state does not fully define the state of the system—which is likely given that the internal neural as well as environmental variables are ignored—it is easy to imagine that this mapping will not be right-unique and thus, mathematically speaking, seize to be a function. The SM environment is an even more general relation, a superset of multiple forward models. In discrete terms, it would have the form $R(m, s_{t+1} - s_t)$.

[10] Note that the Section's name is "Extracting a body schema from raw sensorimotor data", illustrating the confusion of terms.

[11] The particular details differ though. First, due to the dimensionality of the sensorimotor space, we studied relationships between pairs of variables only. On the other hand, as opposed to SM environment, we included sensory-sensory pairs as well. In addition, we applied a particular information theoretic measure, transfer entropy, which allowed us to assess the amount of directed information transfer between individual variables. In this way, information was compressed and salient relationships could be discovered, but at the same time, it did not contain all the information present in the original data.

or for turning [26] and thus correspond to patterns that are functionally relevant for the robot and even carry a normative aspect. Thus, our findings about the sensorimotor space using the gaits (results shown in [51]) can be interpreted as studying the SM coordination or even SM strategy of the quadruped robot. In the "adaptive walking" study, a similar repertoire of coordinated gaits was used. While exercising these in different environments, the robot was taking a record of all the combinations of sensory and motor variables (discretized and compressed into features over 2 second intervals). A reward associated with every sensorimotor state was stored and later used to inform action selection. Thus, the items in the associative memory constitute discrete slices that witness and at the same time influence the robot's SM strategies. Buhrmann et al. [8] also highlight how the space of possible sensorimotor trajectories—in the original sensorimotor space—is narrowed down as one goes from SM environment to SM coordination. In our example, we quantified the overall compression of a theoretical full sensorimotor state space as a result of embodiment and the action selection (internal dynamics) in Section 3.3.

5.2 Clarification

Let us now try to directly compare the "cognitive concepts" that we discussed above in terms of their mathematical formulation, representational nature, and site—where they are located. With the help of the analysis that follows, we will fill up Table 1.

- **Mathematical formulation.** As we have argued, a body schema is a very loosely defined notion and to talk about a mathematical formulation is out of question. A forward model, on the other hand, can be defined precisely as a function. SMCs were also defined rather loosely, but acquired a concrete articulation in dynamical systems terms in [8].
- **Representational nature.** The term "body schema" is usually equated with a body representation. It thus seems to imply a "representationalist" view of the mind. Alsmith & de Vignemont discuss this theme in detail in [1]. A forward model is simply a function on motor and sensory variables, which is per se neutral with respect to the "representationalist dispute". Of course, representational nature can be ascribed to it if one posits that this mapping is stored in the brain and "stands in" for some extraneural states of affairs, as done by Clark & Grush [9], for example. The position of SMCT[12] (in its original formulation [39]) on representations was not clear—for sure it was detailed, pictorial representations, "mirrors of the world states“, that SMCT was arguing against. Buhrmann et al. [8] in their definition and treatment argue clearly against a representationalist interpretation and show that the SMCs—as trajectories in the sensorimotor space—are emergent from the dynamics of the body, brain, and environment (similarly to the dynamic directives proposed in [16]).

[12] Sensorimotor Contingency Theory

– **Site.** A body schema is usually thought to reside in the brain—even if in a highly distributed manner, encompassing for example area SI, area 5 in the parietal lobe, and premotor cortex [21]. The existence of forward models in the brain is also supported by extensive literature, in particular on the cerebellum (e.g., [31, 11]). SMCs are a result of the joint dynamics of the brain, body, and environment; an analysis of the simple agent in [8] reveals that "there is nothing in the internal dynamics of the agent's "brain" that represents the SMCs that are being enacted or the non-actualized sensorimotor regularities that still have a dynamical influence." Yet, some neural vehicles that support the SMCs on the part of the brain seem inevitable and are expected in various brain areas—visual SMCs are discussed in [39]; Engel et al. [16] discuss the role of premotor circuits, for example.

Table 1. Properties of different minimally cognitive concepts.

	Body schema	Forward model	SMCs (according to [8])
Mathematical description	N.A.	Function	Trajectory in S-M space
Representational nature	Yes	Neutral	No
Site	Brain	Brain	Brain-body-environment

6 Robots as Cognitive Science Tools: Are There Intrinsic Limitations?

The methodology adopted in this work was a synthetic one [44]. That is, we built and then studied the behavior of artifacts. As can be seen in Fig. 7, the area spanned by synthetic sciences can be further subdivided into (1) the intersection with empirical sciences—synthetic modeling, (2) the middle area concerned with general principles, (3) the intersection with the application domain in the form of prototypes of new technology.

The scenarios presented here (Section 3) were inspired by skills that were observed in lower animals and serve as instances of the simplest behaviors that we would consider cognitive. Yet, do the case studies presented in this work qualify as synthetic modeling, i.e. as models of biological cognition too? Given that we do not treat cognition as an exclusively biological phenomenon, this possibility is open. However, the parallel between biological cognitive agents and the artificial ones remained on an abstract level—we did not relate directly to any empirical data from the animal kingdom. Along the lines of the critical account in [62], one could argue that this is an example of the "animat" approach to modeling cognition and that more direct parallels to concrete instances of cognitive phenomena in biology are desirable. Several proposals in this direction are put forth in [42]. This would be one possible direction of future work (sketched in Section 8.3.3 in [28]).

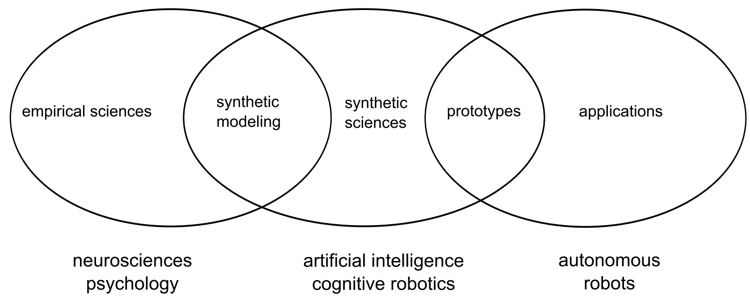

Fig. 7. Overview of approaches to the study of cognition. The figure and caption are adapted from [44] to the study of cognition rather than intelligence. On the left, we have the empirical sciences like neurosciences and psychology that mostly follow an analytic approach. In the center, we have the synthetic ones like AI and cognitive robotics which can either model natural agents (this is called synthetic modeling—the intersection with empirical sciences) or alternatively can simply explore issues in the study of cognition without necessarily being concerned with natural systems. This activity may give rise to prototypes and eventually to full industrial applications, such as autonomous robots.

6.1 The Difficulty of Modeling without Representations

Interestingly, we find that the ease of the synthetic modeling endeavor will depend on the cognitive science paradigm that one is following. Under cognitivism / GOFAI, the body and interaction with environment was of marginal importance, so the robots were not necessary in the first place. Moreover, the representations—models of the world—were often symbolic and directly corresponded to objects in the world (in their designers' eyes). The quality and functioning of the cognitive layer was thus easy to assess. This was obstructed slightly under connectionism, as the models became less transparent due to their subsymbolic nature. Embodied cognitive science then brought about the necessity for considering whole brain-body-environment systems. However, even within embodied cognitive science, the different viewpoints that we have outlined in Section 4 impact on the research methodology. First of all, the move away from veridical to action-oriented or context-dependent representations means that the quality of the internal control structures cannot be assessed by a direct comparison with some objects in the world anymore. The viewpoints that reject representations altogether go even further in this and imply that no answers will be found in the control structure alone [5, 8]. This is analogous to the situation in neurosciences where Engel et al. [16] propose to replace techniques studying neural responses to passive stimuli by studying subjects actively interacting with their surroundings, which brings about many practical difficulties.

6.2 Enactive Robots Subject to Precarious Conditions?

The enactive viewpoint can be taken even further: Di Paolo [13] points out that in order to fully understand cognition in its entirety, embedding the agent in a closed-loop sensorimotor interaction with the environment is necessary, yet may not be sufficient in order to induce important properties of biological agents such as intentional agency. In other words, one should not only study instances of individual closed sensorimotor loops as models of analogous loops in biological agents—that would be the recommendation of Webb [62]—but one should also try to endow the models (robots in this case) with similar properties and constraints that biological organisms are facing. In particular, it has been argued that life and cognition are tightly interconnected [34, 56] and a particular organization of living systems—which can be characterized by autopoiesis [34] or metabolism for example—is crucial for the agent to truly acquire the meaning in its interaction with the world. While these requirements are very hard to satisfy with the artificial systems of today, Di Paolo [13] proposes a way out: robots need not metabolize, but they should be subject to precarious conditions. That is, the success of a particular instantiation of sensorimotor loops or neural vehicles in the agent is to be measured against some viability criterion that is intrinsic to the organization of the agent. The control structure may evolve over time, but the viability constraint needs to be satisfied, otherwise the agent "dies". The unfortunate implication, however, is that research along these lines will hardly fit into the full synthetic methodology scheme (Fig. 7) anymore, since machines whose functioning is not deducible from their control structure and that cannot be given tasks will not easily find their way to application scenarios in industry. On the other hand, this approach may give rise to truly autonomous robots.

7 Conclusion

We focused on autonomous cognitive development and engaged robots in a number of scenarios that can be seen as a developmental pathway from reactive to minimally cognitive behavior. We have experimented with different control architectures and assessed their performance in different tasks. We have also analyzed the nature of these control architectures from the point of view of different cognitive science paradigms. We found that our case studies lend themselves easily to interpretations along the lines of "grounded representation" and internal simulation/emulation theories [4, 9, 22]. On the other hand, if one looks into the details, they are much less compatible with the non- (or anti-) representational or enactive perspectives [58, 54].

The minimally cognitive "building blocks" or notions were also subject to investigations in our case studies. Our results and analysis contributed to a conceptual clarification here. Interestingly, only a forward model seems to be a useful building block that can be deployed in the control structures of robots and serve different purposes—a kind of useful "brain motif" [53] perhaps. A body schema is at the moment an "umbrella term" for a multitude of body representations that can be used for action. This notion is, however, far from a formulation that could be "deployed" in a control architecture. Similarly, SMCs do not constitute a building block either; instead, at the moment, they are rather a descriptive concept, which may prove useful in the analysis of natural and artificial cognitive systems.

Finally, we have evaluated the potential of robots as modeling tools for cognitive science and the implications of this way of modeling regarding the choice of cognitive science paradigm. Adopting an embodied, yet representation-based view, is a convenient choice that creates bridges between the research in psychology, neuroscience and robotics (as elaborated recently under the "grounded cognition" umbrella in [42]). In line with the synthetic approach and a functionalist stance, a particular cognitive architecture may serve as a model of certain parts of the brain and at the same time provide an interesting tool for autonomous robotics, for example. Still, it remains to be shown if human-like levels of complexity can be attained. On the other hand, truly enactive robots seem to be much harder to realize. Models that are compatible with this view are to date of minimal complexity and bear no application potential. From a designer's perspective, achieving an appropriate "shaping of dynamical tendencies that channel appropriate actions on the basis of past experience and in accordance with goals" [8] seems to be much harder than adopting the representationalist stance and tuning a world model of one form or another. Therefore, synthetic enactive approaches in robotics still need to demonstrate their scalability and potential.

Acknowledgement

This work was supported by the project 'Extending sensorimotor contingencies to cognition - eSMCs', FP7-ICT-270212, `esmcs.eu`, and by the Swiss National Science Foundation Fellowship PBZHP2-147259. I would like to thank the collaborators involved in the case studies reviewed in this work, Rolf Pfeifer for continuous support and guidance during this research, and to Giovanni Pezzulo, Alexander Maye, and the two anonymous reviewers for helpful and constructive comments on the manuscript.

References

1. Adrian John Tetteh Alsmith and Frédérique De Vignemont. Embodying the mind and representing the body. *Review of Philosophy and Psychology*, 3(1):1–13, 2012.
2. R. C. Arkin. *Behavior-based Robotics.* MIT Press, 1998.
3. M. Asada, K. Hosoda, Y. Kuniyoshi, H. Ishiguro, T. Inui, Y. Yoshikawa, M. Ogino, and C. Yoshida. Cognitive developmental robotics: a survey. *IEEE Transactions on Autonomous Mental Development*, 1(1):12–34, 2009.
4. L. Barsalou. Grounded cognition. *Annual Review of Psychology*, 59:617–645, 2008.
5. R.D. Beer. The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11:209–243, 2003.
6. R. Blickhan, A. Seyfarth, H. Geyer, S. Grimmer, H. Wagner, and M. Guenther. Intelligence by mechanics. *Phil. Trans. R. Soc. Lond. A*, 365:199–220, 2007.
7. R. A. Brooks. A robot that walks: Emergent behaviors from a carefully evolved network. *Neural Computation*, 1:153–162, 1989.

8. T. Buhrmann, E. Di Paolo, and X. Barandiaran. A dynamical systems account of sensorimotor contingencies. *Front. Psychol.*, 4(285), 2013.
9. Andy Clark and Rick Grush. Towards cognitive robotics. *Adaptive Behaviour*, 7(1):5–16, 1999.
10. T. Collett. Do toads plan routes? a study of the detour behavior of bufo viridis. *J. Comparative Physiology*, 146:261–271, 1982.
11. P.R. Davidson and D. M. Wolpert. Widespread access to predictive models in the motor system: a short review. *Journal of Neural Engineering*, 2:313–319, 2005.
12. F. de Vignemont. Body schema and body image - pros and cons. *Neuropsychologia*, 48(3):669–680, 2010.
13. E. Di Paolo. Robotics inspired in the organism. *Intellectica*, 53-54:129–162, 2010.
14. F.H. Durgin, M. Akagi, C.R. Gallistel, and W. Haiken. The precision of locometry in humans. *Exp Brain Res*, 193:429–436, 2009.
15. A.K. Engel. *Enaction: Towards a new paradigm for cognitive science*, chapter Directive minds: how dynamics shapes cognition, pages 219–243. Cambridge, MA: MIT Press, 2011.
16. Andreas K Engel, Alexander Maye, Martin Kurthen, and Peter König. Where's the action? the pragmatic turn in cognitive science. *Trends in cognitive sciences*, 17(5):202–209, 2013.
17. A.S. Etienne and K.J. Jeffery. Path integration in mammals. *Hippocampus*, 14:180–192, 2004.
18. Anatol Feldman. New insights into actionperception coupling. *Experimental Brain Research*, 194:39–58, 2009.
19. J.A. Fodor. *The language of thought*. Cambridge, MA: Harvard University Press, 1975.
20. Tom Froese. From adaptive behavior to human cognition: a review of enaction. *Adaptive Behavior*, 2012.
21. M.S.A. Graziano and M.M. Botvinick. *Common Mechanisms in Perception and Action: Attention and Performance*, volume XIX, chapter How the brain represents the body: insights from neurophysiology and psychology, pages 136–157. Oxford Univ Press, 2002.
22. Rick Grush. The emulation theory of representation - motor control, imagery, and perception. *Behavioral and Brain Sciences*, 27:377–442, 2004.
23. G. Hesslow. Conscious thought as simulation of behaviour and perception. *Trends in Cognitive Sciences*, 6:242–247, 2002.
24. M. Hoffmann, H. Marques, A. Hernandez Arieta, H. Sumioka, M. Lungarella, and R. Pfeifer. Body schema in robotics: a review. *IEEE Trans. Auton. Mental Develop.*, 2 (4):304–324, 2010.
25. M. Hoffmann and R. Pfeifer. *The Implications of Embodiment: Cognition and Communication*, chapter The implications of embodiment for behavior and cognition: animal and robotic case studies, pages 31–58. Exeter: Imprint Academic, 2011.
26. M. Hoffmann, N. Schmidt, K. Nakajima, F. Iida, and R. Pfeifer. Perception, motor learning, and speed adaptation exploiting body dynamics: case studies in a quadruped robot. In *Proc. Int. Symposium on Adaptive Motion in Animals and Machines (AMAM)*, 2011.
27. M. Hoffmann, N. Schmidt, R. Pfeifer, A.K. Engel, and A. Maye. Using sensorimotor contingencies for terrain discrimination and adaptive walking behavior in the quadruped robot puppy. In T. Ziemke, C. Balkenius, and J. Hallam, editors, *From animals to animats 12: Proc. Int. Conf. Simulation of Adaptive Behaviour (SAB)*, volume 7426 of *LNAI*, pages 54–64. Springer, 2012.

28. Matej Hoffmann. *From locomotion to cognition*. PhD thesis, University of Zurich, 2012.
29. N. P. Holmes and C. Spence. The body schema and the multisensory representation(s) of peripersonal space. *Cogn Process*, 5(2):94–105, Jun 2004.
30. Daniel Hutto. Exorcising action oriented representations: ridding cognitive science of its nazgûl. *Adaptive Behavior*, 2013.
31. M. Kawato. Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, 9:718–727, 1999.
32. M. Lungarella, G. Metta, R. Pfeifer, and G. Sandini. Developmental robotics: a survey. *Connection Science*, 15(4):151–190, 2004.
33. M. Lungarella and O. Sporns. Mapping information flow in sensorimotor networks. *PLoS Comput Biol*, 2:1301–12, 2006.
34. H. Maturana and F. Varela. *Autopoiesis and cognition: the realization of the living*. Dordrecht: D. Reidel Publishing, 1980.
35. A. Maye and A.K. Engel. A discrete computational model of sensorimotor contingencies for object perception and control of behavior. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, pages 3810 – 3815, 2011.
36. T. McGeer. Passive dynamic walking. *The International Journal of Robotics Research*, 9(2):62–82, 1990.
37. G. Metta, L. Natale, F. Nori, G. Sandini, D. Vernon, L. Fadiga, C. von Hofsten, K. Rosander, M. Lopes, J. Santos-Victor, A. Bernardino, and L. Montesano. The icub humanoid robot: An open-systems platform for research in cognitive development. *Neural Networks*, 23(8-9):1125–1134, 2010.
38. H.P. Moravec. The stanford cart and the cmu rover. *Proceedings of the IEEE*, 71(7):872 – 884, 1983.
39. J. K. O'Regan and A. Noe. A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24:939–1031, 2001.
40. N. Oses, M. Hoffmann, and R. A. Koene. Embodied moving-target seeking with prediction and planning. In E. Corchado, M.G. Romay, and A.M. Savio, editors, *Proceeding Hybrid Artificial Intelligence Systems (HAIS), San Sebastian, Spain, Part II*, volume 6077/2010 of *LNCS*, pages 478–485. Springer, 2010.
41. G. Pezzulo. *50 Years of AI*, chapter Anticipation and Future-Oriented Capabilities in Natural and Artificial Cognition, pages 258–271. Springer LNAI 4850, 2007.
42. Giovanni Pezzulo, Lawrence W Barsalou, Angelo Cangelosi, Martin H Fischer, Ken McRae, and Michael J Spivey. The mechanics of embodiment: a dialog on embodiment and computational modeling. *Frontiers in psychology*, 2, 2011.
43. R. Pfeifer and C. Scheier. Sensory-motor coordination: The metaphor and beyond. *Robotics and Autonomous Systems*, 20:157–178, 1997.
44. R. Pfeifer and C. Scheier. *Understanding intelligence*. MIT Press Cambridge, MA, USA, 2001.
45. Rolf Pfeifer and Josh C. Bongard. *How the body shapes the way we think: a new view of intelligence*. MIT Press, Cambridge, MA, 2007. http://mitpress.mit.edu/catalog/item/default.asp?ttype=2&tid=11003.
46. C.G. Prince, editor. *J. Infant and Child Development: Special issue: Developmental robotics*. Wiley Interscience, 2008.
47. Z. Pylyshyn. *Computation and cognition: Toward a foundation for cognitive science*. Cambridge, MA: MIT Press, 1984.
48. M. Reinstein and M. Hoffmann. Dead reckoning in a dynamic quadruped robot: Inertial navigation system aided by a legged odometer. In *IEEE Int. Conf. Robotics and Automation (ICRA)*, pages 617–624, 2011.

49. M. Reinstein and M. Hoffmann. Dead reckoning in a dynamic quadruped robot based on multimodal proprioceptive information. *Robotics, IEEE Transactions on*, 29(2):563–571, 2013.
50. P. Rochat. Self-perception and action in infancy. *Exp Brain Res*, 123:102–109, 1998.
51. N. Schmidt, M. Hoffmann, K. Nakajima, and R. Pfeifer. Bootstrapping perception using information theory: Case studies in a quadruped robot running on different grounds. *Advances in Complex Systems J.*, 16(2-3), 2013.
52. L. Schomaker. Anticipation in cybernetic systems: a case against mindless anti-representationalism. In *Proc. Int. Conf. Systems, Man and Cybernetics*, volume 2, pages 2037 – 2045, 2004.
53. Olaf Sporns and Rolf Kötter. Motifs in brain networks. *PLoS biology*, 2(11):e369, 2004.
54. J. Stewart, O. Gapenne, and E. Di Paolo, editors. *Enaction: Toward a new paradigm for cognitive science*. Cambridge, MA: MIT Press, 2010.
55. E. Thelen and L.B. Smith. *A Dynamic systems approach to the development of cognition and action*. MIT Press, 1994.
56. E. Thompson. *Mind in life: Biology, phenomenology, and the sciences of mind*. Cambridge, MA: MIT Press., 2007.
57. Sebastian Thrun et al. Stanley: The robot that won the darpa grand challenge: Research articles. *J. Robot. Syst.*, 23:661–692, September 2006.
58. F. Varela, E. Thompson, and Rosch. E. *The embodied mind: cognitive science and human experience*. Cambridge, MA: MIT Press., 1991.
59. D. Vernon, G. Metta, and G. Sandini. *Embodied Cognitive Systems*, chapter Embodiment in Cognitive Systems: on the Mutual Dependence of Cognition and Robotics. Institution of Engineering and Technology (IET), UK, 2010.
60. G. W. Walter. *The living brain*. Norton & Co., New York, 1953.
61. B. Webb. Neural mechanisms for prediction: do insects have forward models? *Trends in Neurosciences*, 27:278–282, 2004.
62. B. Webb. Animals versus animats: or why not model the real iguana? *Adaptive Behavior*, 17:269–286, 2009.
63. Margaret Wilson. Six views of embodied cognition. *Psychonomic bulletin & review*, 9(4):625–636, 2002.
64. M. Wittlinger, R. Wehner, and H. Wolf. The ant odometer: stepping on stilts and stumps. *Science*, 312(5782):1965–1967, 2006.
65. Y.P. Yvanenko, N. Dominici, E. Daprati, D. Nico, G. Cappellini, and F. Lacquaniti. Locomotor body scheme. *Human Movement Science*, 30 (2):341–51, 2011.