

Cognitive Agents – A Procedural Perspective relying on the Predictability of Object-Action-Complexes (OACs)

F. Wörgötter^{1,‡}, A. Agostini², N. Krüger³, N. Shylo¹ and B. Porr⁴

- 1) Bernstein Center for Computational Neuroscience, University of Göttingen, Göttingen, Germany, (worgott@bccn-goettingen.de)
- 2) Institut de Robotica i Informatica Industrial (CSIC-UPC), Barcelona, Spain (agostini@iri.upc.edu)
- 3) Cognitive Vision Lab, The Maersk Mc-Kinney Moller Institute University of Southern Denmark, Odense, Denmark, (norbert@mmmi.sdu.dk)
- 4) Dept. of Electrical and Electronic Engineering, University of Glasgow, Glasgow, Scotland, UK (porr@elec.gla.ac.uk)

Abstract: Embodied cognition suggests that complex cognitive traits can only arise when agents have a body situated in the world. The aspects of embodiment and situatedness are being discussed here from the perspective of linear systems theory. This perspective treats bodies as dynamic, temporally variable entities, which can be extended (or curtailed) at their boundaries. We show how acting agents can, for example, actively extend their body for some time by incorporating *predictably* behaving parts of the world and how this affects the transfer functions. We suggest that primates have mastered this to a large degree increasingly splitting their world into predictable and unpredictable entities. We argue that temporary body extension may have been instrumental in paving the route for the development higher cognitive complexity as it is reliably widening the cause-effect horizon about the actions of the agent. A first robot experiment is sketched to support these ideas.

We continue discussing the concept of Object-Action Complexes (OACs) introduced by the European PACO-PLUS consortium to emphasize the notion that for a cognitive agent objects and actions are inseparably intertwined. In another robot experiment we devise a semi-supervised procedure using the OAC-concept to demonstrate how an agent can acquire knowledge about its world. Here the notion of *predicting changes* fundamentally underlies the implemented procedure and we try to show how this concept can be used to improve the robot's inner model and behaviour.

Hence, in this article we have tried to show how *predictability* can be used to augment the agent's body and to acquire knowledge about the external world, possibly leading to more advanced cognitive traits.

‡ To whom correspondence should be addressed at:
Bernstein Center for Computational Neuroscience
University of Göttingen
Bunsenstr. 10
D-37073 Göttingen
Germany

Phone: 0049 551 5176528
Email: worgott@bccn-goettingen.de

Table of Contents

1	Introduction:.....	2
2	On Embodiment.....	4
2.1.	Critical Assessment of the Embodiment Definition	7
3	Object-Action Complexes (OACs).....	7
4	Route to Cognition – Temporary Bodily Integration.....	8
4.1.	Robot Experiment – Temporary Embodiment.....	11
4.2.	Some speculations.....	13
5	Extending the OAC Concept	13
5.1.	Robot Experiments – Discovery by Doing	15
6	Conclusion	18
7	Acknowledgements.....	19
8	References.....	19
9	Appendix A: Body Extension based on Rigid Body Motion.....	23
10	Appendix B: Discovery by Doing.....	25

1 Introduction:

During the last years the European Union has invested more than 100 million Euros into the field of Cognitive Robotics and adjacent fields subsumable under bio-inspired advanced robotics. As the development of such programs rests on the ever growing scientific community in these fields, this is indicative of the fact that “a lot of people believe in it”. It appears, that machines with truly intelligent, cognitive features¹ have now become within reach of research and development. This may have been largely due to the emergence of “Embodied Cognition” (EC) as a possible theoretical foundation for such R&D activities (Lakoff and Johnson, 1999; Brooks, 1999; Todes, 2001; Varela et al., Rosch, 1991). Summarized in one sentence EC assumes that only machines with some kind of a body, which allows direct interactions with the world, hence, which situates these machines in their world, will be able to develop advanced (cognitive) traits (Chiel and Beer, 1997; Pfeifer and Scheier, 1999; Steels and Brooks, 1995; Clancey, 1997; Clark, 1999; Todes, 2001; Riegler, 2002). This notion was much influenced by Rodney Brooks, who was one of the first to explicitly state these ideas in the context of robotics work (Brooks, 1986). Embodied Cognition is thus different from what has been called *good old-fashioned AI* (GOFAI), which in its extreme form supports a Cartesian attitude, treating the mind as an entity independent of and, thus, not requiring, the body (see Anderson, 2003 for a comparison between the Cartesian viewpoint and EC). This article does not intend to enter into the controversy between GOFAI and embodied cognition (first pointed out explicitly by Dreyfus, 1972; see also Brooks, 1999). For our purposes it suffices to just illuminate a little bit the germination process of EC, which has to a large degree been triggered by the notion that after all GOFAI-systems have not really become intelligent (see Brooks, 1999 for a discussion). A wealth of possible problems has been put forward for explaining this. Most influential was here the discussion of the symbol grounding

¹ The term “cognition” is exceedingly ill-defined and no common agreement exist about how cognitive is cognitive. After all, also ants can build houses. We will use the term also in a wider sense but always in conjunction with human cognitive traits. Furthermore, this article is largely devoted to the question, what could be a path towards cognition. Hence we are concerned with processes and not so much with their final outcome. **Cognitive complexity is, thus, as we see it, a continuum.**

problem (Searle, 1980; Harnad, 1990) and the frame problem of AI (McCarthy and Hayes, 1969; Dennett, 1984) as this had prepared the ground for the germs of EC. Meanwhile a large number of articles have appeared discussing, often in a controversial way, which conditions are necessary for embodied cognition. A very nice summary of this is given by Wilson (2002). She wraps up six of the most common claims on cognition found in the literature that to a large degree ask the question about the interaction between agent and world, which is a central topic also of our paper.

However, one aspect of the EC-discussion is quite puzzling in general. Most of the discussion revolves around *necessary conditions* for a cognitive agent. Little is said about what would be sufficient to drive cognitive development. Necessary conditions do not specify any on-line procedure, any ontogenetic developmental process, or any phylogenetic evolutionary mechanism that could actually drive the development of cognition. Hence, from a robotics perspective, necessary conditions are only half of the game. If you cannot show (or at least suggest) a process that leads to the germination of “something cognitive”, not much has been achieved along those lines.

One possible way out of this dilemma was the idea to let robots develop similar to human infants, leading to the growing field of “*developmental robotics*” (Weng et al., 2001; Lungarella et al., 2003). For this idea, we, humans, are the proof of concept. Hence: build a robot, make it similar to a human, endow it with enough sensor-motor complexity, and with a set of useful learning algorithms and let this agent develop and learn in interaction with its world and other agents (usually its designers) and you will see the emergence of cognition. This can be done with real robots, different from the field of “*evolutionary robotics*”, (Nolfi and Floreano, 2000) which attempts the same goals but must almost exclusively rely on simulations, as physical robots cannot have offspring and mutate. Both fields have their successes and increasingly complex behaviour is observed in such agents, which may some day be (or look) cognitive.

What remains frustrating about these approaches is that self-organization might indeed lead to cognition (future will show!), but, we are probably none the wiser as it is exceedingly difficult and many times totally impossible to gain a deeper understanding about the final (developed or evolved) system, let alone about the dynamic processes that have led to it².

While developmental and/or evolutionary robotics may indeed be a way forward, we would nonetheless suggest devoting more effort to the denomination, the theoretical understanding, and the technical implementation of possible *sufficient conditions* for cognition. One key question is: Is it possible to specify some processes that may in a theoretically grounded way lead the way towards cognition in machines? To this end we would like to adopt a systems theoretical perspective on agents and their world (Ashby, 1952; McFarland, 1989; Walter, 1953), which has the advantage that its cybernetic ideology is already “very procedural” as such. In doing so some aspects of biological agents (animals and their nervous system) will be discussed, which appear to be relevant in this context.

² Think about self-organization of neural networks as an example. Many algorithms exist for this and a wide variety of problems can now be solved by ANNs. On the other hand the theory of ANNs is mostly only developed for linear systems and it is very hard to understand more difficult ANNs in an analytic way.

This article is structured as follows. In general we will present several different results and ideas on the questions of embodiment, situatedness and cognition. The core thread which links them is the aspect of *Predictability* and procedures which involve predictability around which these ideas evolve. First (section 2) we would like to provide an improved systems theoretical perspective on embodiment relying on linear systems theory. By this we will define in a more rigorous way what is the body of an agent against “the world outside” (Porr and Wörgötter, 2005). This specification has direct implications on understanding the interactions between agent and world from which cognition might arise. In the next section (3) we will introduce so-called “object-action complexes” (OACs) as possible structural entities relevant for cognition³ (Hommel et al., 2001). Then (section 4) we will suggest a process by which an agent can extend its body-image and show robot experiments for this (4.1) arguing that this might be helping the agent to develop cognitive traits (4.2). Next (section 5) we will extend the OAC concept asking which aspects of objects and actions are relevant for an agent (animal) and define “Change”, “Repeatability” and “Predictability”, falling back on observations from the neurosciences. We will implement these aspects in another simple robot experiment (5.1) showing how a procedure can be devised by which a machine can discover parts of its world. Finally in section 6 we will conclude this article with a discussion.

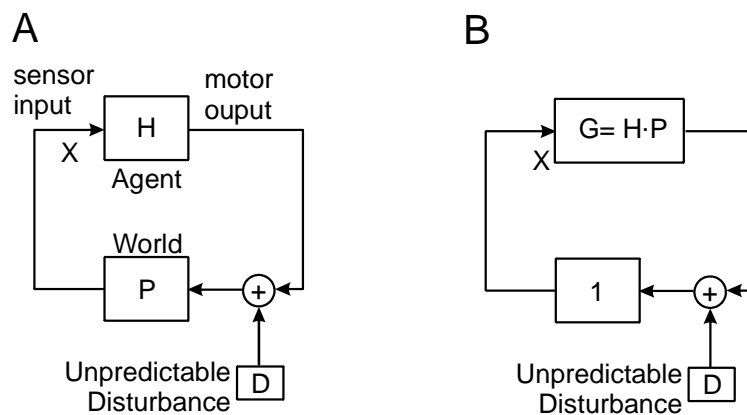


Fig. 1) Applying linear systems theory to define agent and world.

2 On Embodiment

In 2005 we had tried to provide a systems theoretical description of embodiment (Porr and Wörgötter, 2005) from the viewpoint of a constructivist (von Foerster, 1960; Maturana and Varela, 1980; von Glasersfeld, 1996). This perspective shall also be adopted here, because during phylogeny functional traits can only have developed by animals interacting with their environment. This situation is depicted by the simple diagram in Fig. 1A, where H describes the transfer function of the agent and P the transfer function of the world. Sensor inputs arriving at the agent will through H be transformed into motor outputs, while those will – in turn – be transformed into new sensor inputs for the agent through P, the transfer function of the world. For example the lifting of an object will lead to a changed visual sensation (the object moves), where this sensation is different depending on if you have lifted an object submerged in water as compared to air (different refraction index leads to different P). The

³ Note, the concept of OACs is still to some degree “emerging” and being controversially discussed (Geib et al., 2006).

question arises, whether this is an adequate description of system and world. In the first place the chosen formalism is general enough to allow describing all linear systems in a very simple way⁴. We believe, however, that this simple form of describing a system captures the necessary points in the below following discussion context of the differences between embodiment, situatedness and the aspect of temporary bodily extension. Hence, we would argue that it is the strength of the proposed formalism that these aspects do not seem to depend on the specific shapes of H and P. When wanting to discuss more complex aspects of dynamic behavior, however, this chosen description level would probably not suffice anymore.

Both transfer functions H and P form a closed loop which has the task to keep the sensor input X at a desired state. The loop has the task to compensate against unpredictable disturbances coming from D. If we assume that the system is linear, then this leads in terms of closed loop control to the transfer function:

$$X = \frac{D}{1 - HP} \quad (1)$$

With an appropriate design of the functions H and P we are able to generate the desired state at X. It can be clearly seen that the product of H and P defines the property of the loop and how they eliminate the disturbance D.

Note, this diagram describes in the most general sense what it means for an agent to be **situated**. The loop from agent to world and back represents in a systems theoretical diagram the notion of situatedness (Thelen and Smith, 1994; Port and van Gelder, 1995; Beer, 2000). The distinction between H and P corresponds to the distinction between “agent” (the agent’s body) and “world”.

At that point we had argued that every part of the world which can be causally changed by the agent’s output could be integrated into the agent (Porr and Wörgötter, 2005). This can be done by absorbing the environmental transfer function P in H arriving at a new transfer function for the agent $G = HP$ (Fig. 1B).

$$X = \frac{D}{1 - G} \quad (2)$$

This operation can only be performed if P is part of the closed loop, hence for any transfer function which is part of the loop; which is the case for all fully predictable aspects of the world. In Fig.1B this leads to the fact that the transfer function of the world vanishes (becomes equal to one). Examples for an (incomplete) bodily integration process are the “forgetting about” a well-fitting prosthesis, which becomes much integrated in the patient’s body or the feeling of a race-car pilot of “becoming united with the machine”. Note, this definition (Fig. 1A) and the process of integrating P into H (Fig. 1B) do not rely on the physical world. In our sense also non-physical agents (internet robots) can be fully situated and embodied as long as they can be described by such system theoretical relations (e.g. Etzioni and Weld, 1995).

⁴ The assumption of agents as linear systems can be debated because non-linearities can exist in the transfer functions of agent and world. However, the linear system approximation allows a very elegant and clear treatment of closed loop systems which focuses on our main point that transfer functions can be absorbed into each other and which allows a simple way of treating closed loop systems. This could be done similarly in the time domain with non-linear transfer functions but these equations would be much harder to interpret and to discuss. This would eventually lead to a loss in clarity and would rather distract from our main points. Using linear transfer functions in the Laplace domain has a long tradition in cybernetics and dates back to Ashby (1956).

Integration of P into H, however cannot happen for the unpredictable “disturbance” D (Porr and Wörgötter, 2003, 2005). Such disturbance can never be integrated into the body of the agent which can be clearly seen in Eq.1 and Eq.2 where D is actually the input variable and the term $1/G$ the actual transfer function of the loop. This also demonstrates that predictability ($1/1-G$) and disturbance (D) are dependent on each other. The closed loop ($1/1-G$) has been created because there is actually an unpredictable event in the environment. If the agent does not interact with the world, no closed-loop process arises, which would, for example, be the case for a piece of (predictable!) dead matter lying uninfluenced next to the agent. We rather need unpredictable aspects (e.g. weather, other agents, large world, etc) in the environment which will force an agent to create a closed loop system, hence “to react” or even actively “interact” with the world. Note, however, **Predictability** is again only a necessary condition for embodiment. Clearly the trajectory of the sun is predictable and so would be a piece of dead matter but, in spite of this, these entities cannot be integrated into your body.

Two *sufficient conditions*, however, can be suggested to complete the definition.

1) Continuity: The new body part should be integrated for a substantial part of the life time of the agent. Hence any alteration to a body will only with time become a manifest part of the body (the body-image) of the agent. Bodies are continuous for some time.

2) Causal Proximity: The predictable entity, for which bodily integration is to be considered, needs to be proximate (many times resulting in a physical attachment) to the currently existing body of the agent⁵. Causal Proximity is not simply spatial proximity⁶. Above we had argued that an input D (disturbance) needs to interact with the loop $1/(1-G)$. This interaction is a causal process and proximity in our sense means that any entity D, which could potentially be integrated into the body, needs to be part of this closed-loop system. If this is true and if this entity remains predictable for a longer time (continuity), then integration becomes possible.

And so we define:

Entities which are fully predictable and causally proximate to the (current) body of an agent can be integrated into the body. This integration will lead to an alteration of the agent’s body if it is continuous relative to the life time of the agent⁷.

⁵ This causal proximity does not have to lead to any (mechanical/spatial) attachment. Think of a WIFI connected system.

⁶ Causal Proximity – as by name – leads to the fact that the current body of the agent will be able to exert causal effects on the newly considered body parts. This relation is transitive. The new body part should also causally affect the old body (if only through a load change, after having screwed on the new robot hand). In fact the new body part could be much larger than the old body. Think of a small robot that is being physically integrated into a big plant – what is body, what is body-part?

⁷ Note, we are here relying on standard logic definition of equivalence defining: U is necessary for R, formally $R \leftarrow U$, with: predictability \leftarrow body, yielding in standard logic: if part of your body then predicable and also $(\neg U \rightarrow \neg R)$, where \neg stands for “not”, yielding if not predictable then not part of your body. For continuity (S) and proximity (T) this yields that these conditions together with predictability will be “necessary and sufficient”. In standard logic we get then equivalence between the terms involved as: $[R \wedge S \wedge T] \leftrightarrow U$, where \wedge stands for the “logical AND” yielding: if [predicable AND continuous AND proximal] then equivalent to body.

Hence, this definition allows understanding the body of an agent in a constructive way (much like “building” a robot). We note that it is not possible to achieve the situation of Fig.1B in full, removing the transfer function of the world entirely. Even if all entities were predictable for an agent, there would still be many that do not fulfill the two sufficient conditions, which would, thus, contribute only to P and never to H.

2.1. Critical Assessment of the Embodiment Definition

It seems that the above definition is now capturing each and everything and one might ask: are there any agents left at all that are – under this definition – non-embodied? More specifically it also seems that the definition of situatedness, by referring to closed-loop interactions between agents and world, is identical or very similar to the definition of embodiment given here.

Indeed, strong similarities do exist! Let us discuss the different possible cases:

1. NOT Situated NOT Embodied
2. Embodied AND Situated
3. Embodied NOT Situated
4. Situated NOT Embodied

Case 1 would refer to open-loop systems that are without body (e.g. pure symbol manipulation systems). This case is only of theoretical interest here as it would represent the archetype of a Cartesian attitude. Case 2 is most common for biological systems and robots and it is the case that we have discussed in Fig. 1. There is however no principle objection why this case should not hold true for more abstract A-life⁸ creatures like internet agents or computer viruses as long as they obey the necessary and sufficient conditions described above within *their* world.

Case 3 can also be immediately understood as this represents all open-loop systems, which do not feed their output(s) back through the environment onto themselves. These systems are not situated. Case 4 makes troubles, it seems. How could a system be situated but not embodied? It is by the sufficient condition of Proximity and Continuity that this situation can be most easily understood and indeed, biological examples of such systems exist, which are swarms of many (embodied) individuals. In a swarm the individuals have only fleeting contact with each other. Hence as a whole the swarm represents a non-embodied (or very weakly embodied) system, which will however indeed influence its environment and also receive feedback from it. Cognitive complexity can arise from such (social) systems, for example the building of termite mounds, etc. In the context of this article we are, however, not concerned with such social, cooperative aspects, which clearly reach out into human societies, too, and will only to a minor degree discuss some implications of swarm-intelligence later.

3 Object-Action Complexes (OACs)

The question arises whether this notion of embodiment might be of any use for our understanding of agents and their cognitive development. This requires considering the processes of agent-world interaction in more detail for which we would like to introduce the concept of object-action complexes (**OACs**). OACs had first been discussed by the European PACO+ Consortium as a possible way to better formalize the requirements for a machine to approach some level of cognitive complexity.

⁸ A-life=artificial life.

OACs are related to state-action transitions e.g. known from machine learning (Sutton and Barto, 1998; Geib et al., 2006). They rest on the suggestion that objects and actions are inseparably intertwined. Starting with Gibson’s notion of affordances (Gibson, 1979): A hollow thing with liquid may suggest drinking. For this we define an OAC formally by $[O \xrightarrow{A} O']$, which says object O suggests action A and transforms under this action into object O' (cup-full to cup-empty) as the final outcome of this action. Note, rigorously one should define the OAC with respect to the **Attributes** (full, empty) of an object that get altered by an action. This should be kept in mind when using the abbreviated $[O \xrightarrow{A} O']$ notation. The notion of OACs, however, goes beyond Gibson and the intertwining of Objects and Actions becomes more evident when considering the role of Actions more closely. While objects may suggest actions, it is often the action(-plan) that defines the *objectness* of a physical thing. This become clear by following example: It is the action of drinking that *makes* this thing _{hollow,full} “a cup” (“a container”, etc.). The decisive influence of the action becomes immediately obvious if you plan to turn the (same!) thing _{solid-bottom} upside down to use it as “a pedestal” for some figurine for your mantelpiece decoration. Hence, the planned and executed action turns a thing with some (required) properties into a meaningful object. Depending on the planned actions, different properties of the same thing (hollow, full vs. solid bottom) become important. Clearly it is a very difficult (cognitive) problem for an agent to find out which properties are important and which are not. We will come to this later.




Fig. 2) Different types of transformations of objects by actions

4 Route to Cognition – Temporary Bodily Integration

In the following we would like to suggest how the above notions might be helpful in defining some processes that could indeed lead to higher behavioural complexity in an agent suggesting a possible route to (higher) cognitive traits.

It has long been known that being able to predict the world, or more specifically to predict the changes induced by the agent in the world⁹, leads to improved fitness of the agent fostering its survival (and reproduction). Furthermore, a whole field has emerged during the last 10 years or so, which tried to explain advanced cognitive properties by so-called “probabilistic models” (Thrun et al., 2005; Chater et al., 2006; see also a special issue in TICS, 2006; on Probabilistic Models of Cognition). These models most often rely on Bayesian inference (Bayes, 1763; Tenenbaum et al., 2006) which is a powerful probabilistic method for making predictions.

⁹ It may make sense to point out that we are talking about Predictability from an agent centred perspective (actions by the agent). Predicting events that happen in the world without the agent’s doing will also improve fitness (“thunder may predict rain”). This refers to temporally related events, which follow each other, such that this correlation can be learned. This is, however, an entirely different type of predictive mechanism not of relevance in the context of this paper.

A frame problem, however, hides here (McCarthy and Hayes, 1969; Dennett, 1984). In a complex world, such as that a robot or human faces, it does not make sense to try to predict each and everything. It is, thus, of interest to analyse a bit more in detail from a procedural perspective what happens when an agent interacts with the world using the OAC concept and bringing it together with some systems theory. This viewpoint will lead to the notion of “*being predictable and, therefore, we can ignore it*” as a powerful method allowing the agent to free mental resources and avoid such a possible frame problem from starters (see “Some Speculations”, below).

Figure 2A shows that during the interaction of an agent with an object normally also attributes of the agent (O_2) will change¹⁰. After all the effectors of an agent are also just physical objects that will be influenced when getting in touch with another object (O_1). For biological agents such contacts are most of the time fleeting and of little duration as indicated in Fig. 2A by the small contact zone of both OACs. An example would be a cat chasing a ball around. A different situation is depicted in Fig. 2B. Here a more permanent contact is established between agent and object established by action A_1 . Such cases also exist for animals (a cat holds a mouse between its fangs). A new object O_3 has this way been formed, however, for most animals follow-up actions are normally very restricted and object manipulation cannot be performed (beyond the eating of the mouse). This is different for humans and in a restricted way for some animals (Hunt, 1996; Povinelli, 2000; Weir et al., 2002). Dexterous manipulation becomes possible by the fact that we can use the newly formed object and move it in a *predictable* way using our hands leading to situation (C) in Fig. 2. This notion is not terribly new as such but some interesting conclusions arise when looking at the situation in Fig. 2C from a systems theoretical viewpoint.

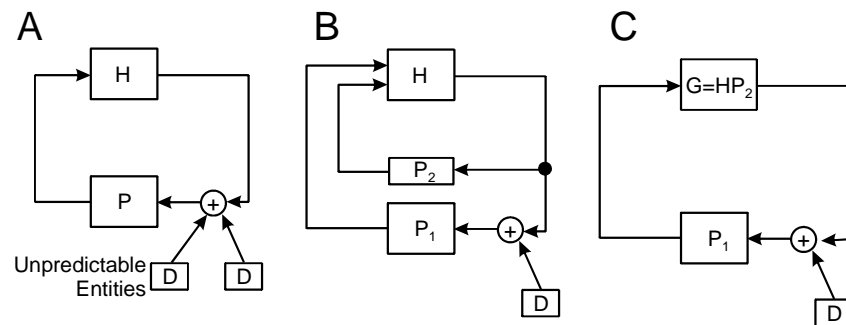


Fig. 3) Systems theoretical representation of temporary embodiment.

Figure 3A depicts a situated agent (human) facing a few disturbances. In the process of grasping an object the human will – if successful – be able to make the grasped object fully (or at least very) predictable to her. Hence, an entity that had been a

¹⁰ According to Petrick and Geib (personal communication) the aspect of an object’s attribute and, thus, the OAC definition as such, needs to be more carefully considered. Think about an open door which affords the action of walking through, by which the door’s attribute (open) will not change. We believe that this does not pose a problem for the definition of OAC as given above, though, as the attribute list of the object, can in principle also contain entries about the relation of the agent with respect to the object. Figure 2A suggests that by an action the agent will also change. Or more specifically the relation of the agent to the object changes (panels B, C). As O_1 and O_2 are symmetrical, one could attach an attribute to one of them (or to both) which describes their relation to each other. Through the performed OAC, this attribute will change.

disturbance (of – say – her visual input space) will first become a predictable entity P_2 (Fig. 3B), where the human will then be able to *temporarily integrate* this entity into her body (Fig. 3C). The remaining aspects P_1 of the world cannot be integrated as they might, for example, be too unpredictable or too far away or from the agent's currently existing body. The idea that humans (and monkeys) indeed perform temporary bodily integration is supported by experimental results that over time cortical receptive fields are extended representing the tip of a stick, which a monkey had to use to obtain food for an prolonged period of time (Obayashi et al., 2000, for a very nice review see Iriki and Sakura, 2008). Hence a long duration, where the processes depicted in Fig. 3 had taken place, has in this case even led to a long-lasting plastic change of the nervous system of this agent (monkey).

The apparently strange notion of temporary bodily integration becomes much more digestible if one thinks of an advanced robot that has grasped a pair of pliers and can handle it now with high precision and dexterity. What would prevent us – the robot's designers – from using a few screws to permanently attach these pliers to the body of the robot this way making the temporary bodily integration a permanent one?

This brings us briefly back to swarms: Here one could argue that (social) contacts formed between individuals would lead also to an augmented body concept by which a swarm can achieve more than any of its members. For slime moulds such contact can indeed be permanent and they can, indeed, form a body in the more traditional sense. Hence, it seems that gradual transitions and different types of temporary bodies do indeed exist. It would be interesting to look at swarms and swarm re-organization also from a systems theoretical perspective, but this would go beyond the scope of this article.

A possible additional complication arises from the fact that the above presented picture seems to suggest that “the world is deterministic”. Instead, usually the change of an attribute takes place on continuous stochastically affected variables. For example the filling level of a cup resulting from the OAC “filling the cup” depends on how good the agent can perform this action. A clumsy agent will produce widely varying outcome. Hence, at first, action outcomes normally need to be measured in continuous space and sometimes also monitored along continuous time. This makes it absolutely necessary to include continuous models for attributes in the process. Second, the reduction of stochastic influences can be performed by model update which can rely on repetition (of the same OAC) and averaging, but possibly also on reinforcement learning or even on supervision (a teacher tells the agent how full the cup should be). Such model update mechanisms on attributes, which have a variance, are important processes and we will deal with this when presenting the experiments of the section “Discovery by Doing” below. However, as such these considerations do not alter the line of system theoretical argumentation presented above.




Fig. 4) Temporary embodiment experiment (for explanation see text).

4.1. Robot Experiment – Temporary Embodiment

In the following we will describe a set of experiments performed with a simple industrial robot (Stäubli, Switzerland) demonstrating how the principle of temporary bodily integration can be implemented in a machine in a simple algorithmical way to provide some support to this idea. A more detailed technical description of this is given in Appendix A.

To this end we assume a few things for our machine to be innate:

- A. A visual representation exists by which a scene can be decomposed into simple 3-D entities, which we call primitives (see Fig. 4A especially also the inset; for technical details see Krüger et al., 2004). Notions of distance (metric) exist therein.
- B. The machine “knows” that coherently moving primitives belong together. This is known as the rigid body motion principle (see Faugeras, 1993) and corresponds to the prominent Gestalt law of Common Fate.
- C. Through this, the robot is able to learn about its own body (gripper). This can be achieved by a purely correlation based learning process where the robot has learned to associate coherent motion in the visual field to the fact that there has been a motor command, which the machine has used to perform a movement. We assume that the process of knowing its own body is basically completed; but that this process keeps on running “in the background” to safeguard against incompleteness and errors in the body representation.

- D. The machine can move its arm and it has also a certain drive to move its arm around (without which nothing would ever happen!)
- E. The machine can push things around by making (visually measured) contact to entities in the scene, which do not belong to the machine. Measurement relies again on the 3-D primitives for which the concept of distance exist.
- F. A grasping reflex can also be performed with some success, triggered by certain geometrical constellations between primitives from the world (Aarno et al., 2007; Kraft et al., 2007). It can feel a successful grasp (haptic sensors) and it knows that it cannot perform another grasp without first letting go. Like babies, it, however, rather likes to hold on to a grasped entity. After some longer time it might however “get bored” and then it releases the object (also similar to small children).
- G. It has an exploration drive by which it will first try to grasp a thing and if this fails (measured by the haptic sensors at the hands) it will try to push it instead. This exploration is triggered by novelty and will start as soon as something new (new primitives) are discovered in the scene. Novelty is a strong mechanism in animals to elicit such exploration as it strongly stimulates the dopaminergic system of the brain (for review see Fellous and Suri, 2003). The implemented drives themselves are similar to basic reflexes, existing in animals, where for example frogs will try to catch any small, dark fast object which suddenly enters their field of vision, or humans will direct their attention and gaze to bright (red) moving objects.

Note most of the built-in prior knowledge (A-G) could be seem as relatively low on the cognitive scale. Grasping or pushing reflexes require some sensor-motor coupling, the Birth of the Object needs the rigid body motion principle (Gestalt Law of Common Fate), and the agent needs to be able to measure correlations (e.g. for the Birth of the Object and to assess its own body structure). Even simple vertebrates like frogs are know to possess these traits. The goal of the following paragraph will be to show that these aspects can lead on to more complex properties like the temporary extension of the body-image of the agent. This new trait is again still not “deeply” cognitive, but – as argued later – the widening of the cause-effect horizon of the agent achieved this way, could well have been one germ for cognition.

Figure 4A shows the body-representation of the robot as viewed by itself. All black-grey¹¹ primitives have been learned earlier (process C) to belong to its body. In the following we will for simplicity use the primitive type “black-gray” like a mental concept to graphically depict if a primitive is deemed to belong to the robot. If an object enters the visual scene the robot will try to grasp it (process G). If unsuccessful it will push it around (process E). This is shown for a not-graspable, upside down, green cup in Fig. 4B, where three movement stages are shown (Fig. 4B1-B3). If a grasp is successful (Fig. 4C), it will move the object (process D) like the spoon in Fig. 4C, where we show seven snapshots of movement stages. At first it will realize that the object is represented by *many* primitives which belong together (process B). This, we had at some point called “Birth of an Object” as it represents a step where the physical “object-ness” of otherwise purely visual entities (the primitives) can be ascertained (Kraft et al., 2007). If the machine does not accidentally drop the object

¹¹ Note, primitives at an edge are always showing two colours, one for the inside, the other for the outside. Here the robot appears dark (black) and the background brighter (grey) leading to a black-grey primitive.

but instead moves it for a longer time it will realize that the movement of these primitives will (albeit in a complicated geometrical way) be related to its own motor actions (process C). As it does not know better it will update its body-image based on this sensor-motor correlation and extend it to now include the coherently and predictively moving object (process C). This is shown in Fig. 4C by the gradual spread of the black-grey primitives along the (originally red-gray) spoon until the whole spoon is being re-coloured. Again we emphasize that this is just a graphical representation of the spreading inclusion of the spoon into the body image of the robot, similar to what happens in monkeys after extended tool use (Iriki and Sakura, 2008). If a new entity will enter the visual field now, process G is triggered again. It feels reluctant to let go (process F) and, thus, another grasp is inhibited (also F), hence processes G,E will lead to a pushing action now (Fig. 4D). As a consequence this agent, based on very primitive processes, begins to perform an interaction between a very simple “tool” that extends its body (until it drops it) and the world.

Figure 4 shows the complete experiment on our robot in a schematic way. Clearly there are many more rather technical details that we had to take care of until the robot actually could do all this (see Appendix A as well as Krüger et al., 2004; Aarno et al., 2007; Kraft et al., 2007; Pugeault, 2008 for details), but the complete sequence as such does not require any other component beyond those (A-G) listed above.

4.2. Some speculations

What might have been the consequences of such a process for early humans (and possibly for nowadays robots)? Clearly, temporary bodily integration spatially extends the body of the agent and creates a totally new situatedness. As it is predictable, the agent does not have to worry about the new entity and it can largely ignore it “as such”. Instead it can now use it to influence the world, thereby vastly enlarging its contact points to the world, which, before such processes came into being, had been limited to his original, non-extended body only. An agent who has realized that through such a process entities of the world can be made predictable might also have the chance to discover that it is also possible to operate with predictable outcome on other objects out in the world and that chains of predictable outcomes can be actively generated (Mendes et al., 2007), even possible transforming and shaping the object to a certain ends (Hunt, 1996; Weir et al., 2002). The discovery of the Law of Cause and Effect (Thorndike, 1911) during evolution of humankind, which appears to be a cornerstone of complex cognition¹², could, thus, well have been bootstrapped by the horizon enlargement of an agent via temporary bodily integration of parts of the world.

5 Extending the OAC Concept

Above we had defined an Object-Action Complex by $[O \rightarrow^A O']$, where we remind the reader that this is an abbreviation, because, rigorously, one should define the OAC with respect to the **Attributes** (full, empty) of an object that get altered by an action. This should be kept in mind when using the abbreviated $[O \rightarrow^A O']$ notation. Based on this we can also define the **Change** as $\Delta O = O - O'$. Sometimes only the absolute value of ΔO is of relevance, where Change is measured at the object’s attributes which change. Also here we need to add a note of caution. The minus sign does not

¹² We note, that also other animals have some concept on the basic Law of Cause and Effect (many mammals, ravens, etc), but usually only of order one. Hence causal chains remain inaccessible to them.

necessarily represent a mathematical subtraction. This would hold only for attributes which can be encoded in a metric way (like the filling level of a cup). More generally, the minus sign is a comparison operator, which compares an attribute before and after an action. For non-metrically encoded attributes (e.g. rank attributes, discrete non-countable attributes, etc.) the respective comparison operation needs to be differently defined. This comparison procedure requires a memory process because we need to remember a (perceptual) Prior (O) and compare it with a Posterior (O'). Similar to the statement we had made above, we can again point out that it is not an easy problem to determine what (which attributes) should be remembered and compared. When repeating an OAC the agent can also assess the **Expected Change**, which is $\langle \Delta O \rangle$, the average Change across trials, together with its standard deviation $\sigma_{\Delta O}$ called **Repeatability**. A small standard deviation would represent a high repeatability of this OAC, because all Changes are similar. Hence, the Expected Change and Repeatability can be considered as an **Inner Object Model** of a certain OAC. Here we need to strongly emphasize that different inner models are also possible. Change may be very relevant in general, but sometimes the absolute outcome value or a normalized one might be of greater importance for measuring the success of a task. Hence, which model to use, depends on the goals and the task of the agent.

But, which attributes are important? When performing a certain OAC many things can change: Filling a cup leads to a full, heavier cup, an emptier, lighter coffee pot a splashing noise and possibly a change of the illumination (because someone else has switched a light on). Normally, through repetition the agent can find out which properties change causally (Thorndike, 1911), hence in way correlated to the OAC (certainly not the illumination), and this way the agent can improve the Expected Change reaching smaller values of $\sigma_{\Delta O}$ with more and more trials¹³. Here we note that we have tacitly assumed that the agent will be able to perform “fairly optimal” actions. Hence, the filling-action as such ought not to introduce additional contingencies which affect Change and Expected Outcome. In reality – say for small children – this is not necessarily the case. Hence such “clumsy” agents need to improve their actions in parallel to updating the Expected Outcome of the corresponding OAC. Thus, **Action Models** *as well as* Inner Object Models need to be updated and improved in parallel. Assuring convergence of such a double-procedure is far from trivial and would require special algorithmic attention.

Storing Expected Change and Deviation of course requires also some kind of memory. A simple way to measure the **Unpredictability** U of an OAC is to calculate $U = \text{abs}(\Delta O - \langle \Delta O \rangle)$, Change for a single trial minus Expected Change¹⁴. If your actions are ok, and your inner object model is complete you should find that actual action outcome and expected outcome match. Hence, this OAC is then highly predictable and U is small. Humans keep on trying when an attempted OAC keeps on leading to an unpredictable outcome. On the other hand, they are getting bored and stop repeating actions for which a predictable outcome is again and again obtained. Boredom¹⁵, hence comes with high Repeatability which corresponds to a small value

¹³ Alternatively a teacher can tell the agent to “pay attention to” a certain change (similar what we often do with our kids), which is a much more efficient procedure as compared to many required repetitions in trial and error learning.

¹⁴ We define Unpredictability as predicting change, not predicting status-quo.

¹⁵ It is equally possible to use these quantities to define other psychologically relevant entities, like Frustration, which would arise if Repeatability does not get better.

of $\sigma_{\Delta O}$. Boredom, however, will not be considered in the example below. One more thing should be reasonably assumed. Agents should not average Change and Predictability over their complete life span, Forgetting – hence limiting the averaging window – will help to remove the influence of early trials on the inner model (which are most of the time much more erroneous than later trials).

The notions above heavily rely on the aspect of Change, Expected Change and (Un-) Predictability. The nervous system of (probably all) animals is highly change sensitive and vastly neglects constant inputs as almost all neurons respond in an adaptive way (“phasic”) bringing their activity levels back to (near) zero if their inputs remain constant. Hence, without over-stating, it seems fair to say that Change is the most relevant aspect of the world for biological agents.

5.1. Robot Experiments – Discovery by Doing

The considerations above suggest that Predictability will allow an agent to extend its body, but it will also allow for “discovering the world” in a reliable way. This has been done in another set of experiments based on the learning of cause-effect rules by the help of a human teacher. More details on the algorithmic procedure are given in Appendix B and in Agostini et al. (2008a,b). The underlying idea of this is that the robot simultaneously learns cause-effects relations while experiencing situations in accordance to goal oriented behaviour. A cause-effect pair represents an OAC.

For the following it is important to clearly state our goals. The procedure described below focuses on how an agent can enter an OAC into its rule-set and how it can statistically update existing OACs, whenever they are being repeated. We are not interested in the different possible ways of how to learn an OAC. Thus, we implemented a largely supervised learning procedure, by which, if necessary, a human teacher instructs actions and explains outcomes. This is a very efficient way of teaching but could be replaced by trial and error learning, reinforcement learning or other learning methods if desired. In the following we will also not discuss how an agent would make a choice (*decision making*). This is a different issue unrelated to the model building for OACs. The same is true for *planning*, which is also not of importance here. Specifically we will focus on a procedure for generating updating OACs based on the concepts of Change, Predictability, etc, developed above.

Table 1) Definitions

ΔO	Change
$\langle \Delta O \rangle$	Average Change
$\sigma_{\Delta O}$	Repeatability
U	Unpredictability
Φ_S	Success Threshold
N	Counter

Let us, therefore, come back to our definition of Change, Expected Outcome, and especially Unpredictability $U = \text{abs}(\Delta O - \langle \Delta O \rangle)$, Change for a single trial minus Expected Change¹⁶. As the definition of U relies only on (cumulative) changes it is independent of the specific OAC performed and we can measure the **Success** of any OAC by holding U against some threshold Φ . If $U < \Phi_S$ the OAC was successful. Clearly, success thresholds should change with praxis. An inexperienced young child building a LEGO toy house will have a different (in this

¹⁶ We note that all these definitions apply to one given OAC. Any other OAC will have values for these entities on its own. This would require an index, which we would like to avoid for easier writing and more clarity.

definition: higher) success threshold than an experienced adult. Especially, when performing an OAC for the first time, its success threshold should be very high, because the agent does not know about any Expected Change. **Failure** occurs for $U > \Phi_S$ and this may lead to **Surprise** and the triggering of a different OAC, possibly for resolving the Surprise, or it may lead to the repetition of the current OAC to check the consistency of the Failure. Thus, Surprise arising from an unpredictable situation can be one strong driving force for the discovery of the world by a machine. Note, Surprise is not the only driving force. Even without Failures (which lead to Surprise) exploration can be triggered by boredom and other psychological variables.

With these definitions it is possible to arrive at an interactive algorithm where an agent performs known OACs and/or learns new ones relying on supervision by a human. The human teacher will basically explain all unknown situations to the agent, if necessary telling the machine to perform a certain action and to pay attention to a certain attribute after perceiving the action outcome. Hence on every first encounter the complete cause-effect relation covered by a new OAC is being supervised. This removes the terrific problem of how to discover cause-effect relations (which could be done by measuring correlations through many repetitions of the same OAC) and makes this learning process highly efficient, similar to the teaching of an attentive human student. Following our goals, the system itself, however, includes the building and updating of inner models which represent each OAC. Hence, the agent will, when repeating an OAC, measure Change ΔO , calculate Unpredictability U , and update Average Change $\langle \Delta O \rangle$ as well as Repeatability $\sigma_{\Delta O}$. On success it will also subsequently lower the success threshold Φ_S . This way successes and failures can be continuously monitored, where the agent will again “call for an explanation” in case of a failure.

The system, which we have implemented, is described in more detail in Appendix B. Here we would like to show only an example (Fig. 5) where the goal is to clear parts of a table by moving bowls out of the way with a robot arm to make it possible to move one bowl (dark) to the far end. See Fig. 7A for a picture of the setup. To this end the agent uses a grid of locations and learns to interpret what it means for a given action that grid locations are occupied or not¹⁷.

The pseudo code below just shows the OAC update subroutine based on the variables above. The currently perceived situation is stored in s_{Prior} which is part of a perception function used to store the perceived situations (see Appendix B). OACs rely in this formalism on consecutively occurring perceptual states (s_{Prior} , s_{Post}), on an action, which relies on a Task coupled to s_{Prior} . The action transforms s_{Prior} into s_{Post} . Furthermore an OAC contains models of expected outcome measured as the Average Change $\langle \Delta O \rangle$ together with the assessment of the Repeatability $\sigma_{\Delta O}$ for the respective OAC. Each OAC carries a Success Threshold Φ_S .

¹⁷ It is interesting to ask, what are actually the objects in such a scenario? Evidently the bowls are not the objects for the OACs, which we are considering here. Rather, it is the (imagined) grid locations, which take the characteristics of an object here. Their attributes are if they are occupied or not. It is one strength of the OAC concept that it also holds in the context of such an example.


```

Subroutine: Execute and update OAC(SPrior, Task)
Execute action (deduced from SPrior, Task or instructed by teacher)
Perceive posterior SPost
Measure Change  $\Delta O$  (Posterior - Prior)
Calculate  $U = \text{abs}(\Delta O - \langle \Delta O \rangle)$  (UNPREDICTABILITY)
  IF not the expected outcome, hence  $U > \Phi_S$  (FAILURE=SURPRISE)
  {
    Teacher explains bad/missing conditions (with indications like "detector A needs to have value B"
    to afford the expected outcome.)
    Correct OAC using teacher explanations (by replacing the value of detector A with
    attribute B in the OAC representation.)
  }
  ELSE (if not failure = success)
    Update OAC statistics: update  $\langle \Delta O \rangle$ ,  $\sigma_{\Delta O}$ , and  $\Phi_S$ 
Return

```

Figure 5A shows a representation of the initial state and the goal in our particular task. It also illustrates the sequence of OACs which have to be executed to reach the goal. The dashed cells indicate a “don’t care situation”, where the agent will not consider if there is a bowl or not at this location. Hence the agent assumes that the occupancy of these particular grid cells does not affect the OAC.

This set of OACs (A) will be fully instructed by the teacher in the first trial of the agent, because the agent has no knowledge whatsoever. During subsequent trials more and more OACs will be learned this way and when an OAC is being repeated it will lead to an update of its statistical properties, Average Change $\langle \Delta O \rangle$, Repeatability $\sigma_{\Delta O}$, and Success Threshold Φ_S . This is depicted in Fig. 5B where OAC2 from (A) is repeated in a much different situation leading to an update of the expected outcome. The agent will now give more trust to the fact that cell 2 should be occupied after having performed this particular OAC (depicted by a darker grey shading of cell 2 at the bottom). The estimate of Average Change will get more stable, Repeatability will improve and Success Threshold will be lowered.

Evidently, OAC4 in panel (A) is still incomplete. With the existing knowledge from trial (A) the agent assumes that the preconditions for this OAC are: cell 10 must be occupied, cell 14 must be free and all other cells do not matter. This is incorrect as shown in Fig. 5C, because cells in between also need to be free to be able to perform the action 10U4. In the situation depicted in panel (C) surprise will arise and the agent will “discover” (by the help of the teacher) another relevant precondition for this OAC.

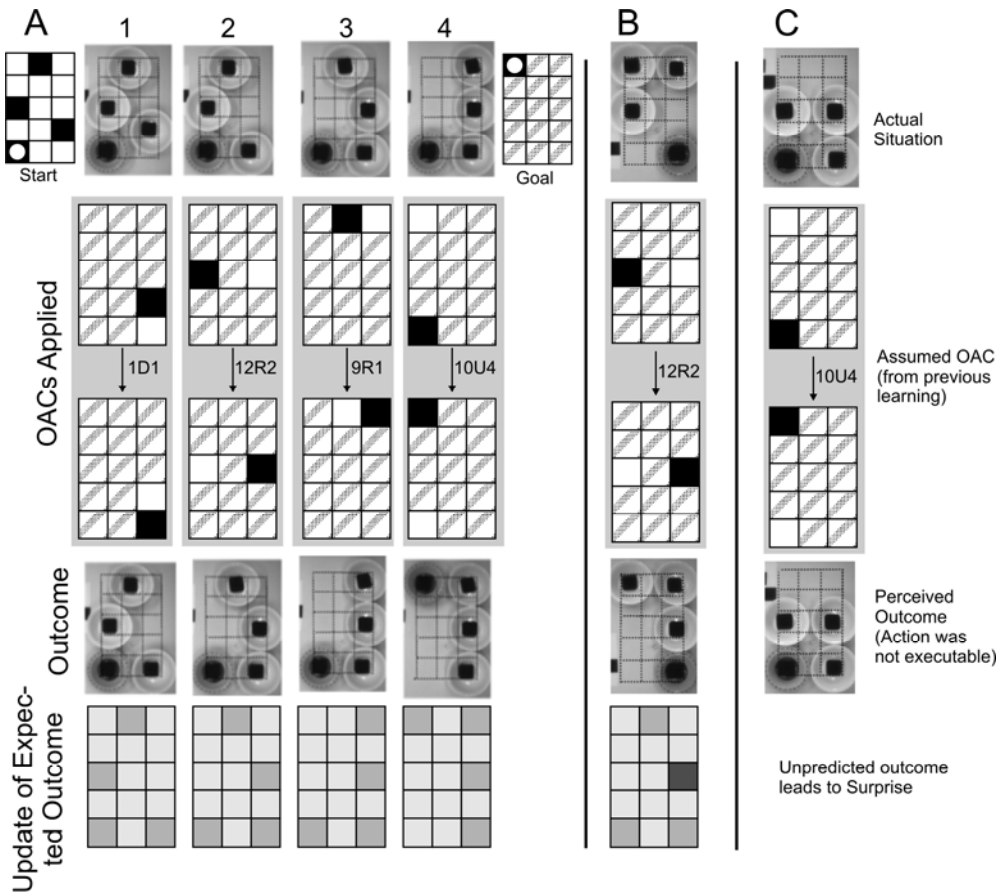


Fig. 5) Experiment with a robot arm for learning from a human to clear parts of a table and to move one object there. For further explanations see text.

6 Conclusion

Many of the notions put forward in this article had been presented at least in parts by other authors, which we have tried to acknowledge along the way. So this paper hopes to contribute to the discussions on embodied cognition by its different (systems theoretical) perspective as well as by the attempt to find an uninterrupted procedure, based on the evaluation of Change and Predictability, towards more cognitive complexity. “Uninterrupted” means here that the same principles have probably been valid for our primordial ancestors and are still applicable for us and our children (which we teach). It is plainly impossible to draw all cross-links to prior work and some aspects have been totally left out, for example, we have totally omitted Piaget’s view (Piaget, 1930).

Thus, in this article we have tried to provide a procedural perspective on embodiment and cognition using ideas from linear systems theory to explain our assumptions. This approach allows disentangling the concept of embodiment from that of situatedness and relies heavily on concepts of “Change” and “Predictability”, which are prevalent in neuronal responses. Specifically, we tried to show what happens to a system when an agent is able to manipulate an object in a predictable way: From a systems theoretical viewpoint this object then becomes temporarily integrated into the agent’s body. By ways of simple robot experiments we have shown that the idea of temporary bodily integration can be consistently represented on a machine using the

principle of rigid body motion (RBM) to integrate entities into the body image of the agent as soon as those entities move coherently and if this happens together with a motor command that the agent has produced. Given preconditions A-G above, this process relies then only on signal correlations and does not need any teacher or other external influence (which could not have been there anyways during evolution).

In the following we had discussed how a process (relying on Change, Predictability, Surprise, etc.) can be formulated by which a robot can (with help) discover the rules of its world. This is achieved by exploration, repetition, surprise and the resolving of surprise, which can either be achieved by a teacher or through trial and error learning. Trial and error learning makes such procedural trees accessible already early in human evolution, where supervision has still not played a big role and learning was almost exclusively by trial and error.

7 Acknowledgements

This work was funded by the European project PACO-PLUS. Thanks are due to Mila Popovic and Nicolas Pugeault for their help with the robot experiment in Fig. 4 and to T. Asfour, C. Geib, B. Hommel, R. Petrick, M. Steedman and the rest of the PACO-PLUS consortium for many stimulating and fruitful discussions. Many thanks also go to S. Wischmann and M. Tamosiunaite for their help with the concepts of this paper and to Tomas Kulvicius for general help.

8 References

D. Aarno, J. Sommerfeld, D. Kragic, N. Pugeault, S. Kalkan, F. Wörgötter, D. Kraft, D. N. Krüger, Early reactive grasping with second order 3D Feature, in: IEEE International Conference on Robotics and Automation (ICRA), Workshop: From features to actions – unifying perspectives in computational and robot vision, 2007.

A. Agostini, E. Celaya, C. Torras, F. Wörgötter, Action rule induction from cause-effect pairs learned through robot-teacher interaction, in: Proceedings of the International Conference on Cognitive Systems, CogSys 2008, Karlsruhe, Germany, 2008a, pp. 213-218.

A. Agostini, E. Celaya, C. Torras, F. Wörgötter, Learning rules from cause-effects explanations, Technical report, Institut de Robòtica i Informàtica Industrial, UPC-CSIC, Barcelona, Spain, 2008b.

M.L. Anderson, Embodied cognition: A field guide, *Artificial Intelligence* 149 (2003) 91-130.

W.R. Ashby, *Design for a Brain*, Chapman and Hall, London, 1952.

W. R. Ashby, *Introduction into Cybernetics*, Methunen, London, 1956.

T. Bayes, *An Essay Towards Solving a Problem in the Doctrine of Chances*, <http://www.stat.ucla.edu/history/essay.pdf>, 1763.

R.D. Beer, Dynamical approaches to cognitive science, *Trends in Cognitive Sciences* 4 (2000) 91-99.

R. Brooks, A robust layered control system for a mobile robot, *Journal of Robotics & Automation* 2 (1986) 14-23.

- R. Brooks, *Cambrian Intelligence: The early History of the New AI*, MIT Press, Cambridge, MA, 1999.
- N. Chater, J.B. Tenenbaum, A. Yuille, Probabilistic models of cognition: Conceptual foundations, *Trends in Cognitive Sciences* 10 (2006) 7:287-291.
- H. Chiel, R. Beer, The brain has a body: Adaptive behaviour emerges from interactions of nervous system, body, and environment, *Trends in Neurosciences* 20 (1997) 553-557.
- W.J. Clancey, *Situated Cognition: On Human Knowledge and Computer Representations*, Cambridge University Press, Cambridge, 1997.
- A. Clark, Embodied, situated and distributed cognition, in: W. Bechtel, G. Graham (Eds.), *A Companion to Cognitive Science*, Basil Blackwell, 1999, pp. 506-517.
- D.C. Dennett, Cognitive wheels: The frame problem of AI, in: C. Hookaway (Eds.), *Minds, Machines and Evolution*, Cambridge University Press, Cambridge, 1984, pp. 129-151.
- H.L. Dreyfus, *What Computers Can't Do: A Critique of Artificial Intelligence*, Harper and Row, New York, 1972.
- O. Etzioni, D.S. Weld, Intelligent agents on the internet: Fact, fiction, and forecast, *IEEE Expert* 4 (10) (1995) 44-49.
- O.D. Faugeras, *Three-dimensional Computer Vision*, MIT Press, 1993.
- J.M. Fellous, R. Suri, The Roles of dopamine, in: *The Handbook of Brain Theory and Neural Networks*, second edition, edited by Arbib, MA, MIT Press, 2003.
- C. Geib, K. Mourao, R. Petrick, M. Pugeault, M. Steedman, N. Krüger, F. Wörgötter, Object action complexes as an interface for planning and robot control, in: *IEEE RAS International Conference on Humanoid Robots*, Genova, December, 4-6, 2006.
- J.J. Gibson, *The Ecological Approach to Visual Perception*, Houghton Mifflin, Boston, 1979.
- S. Harnad, The symbol grounding problem, *Physica D* 42 (1990) 335-346.
- B. Hommel, J. Müsseler, G. Aschersleben, W. Prinz, The theory of event coding (TEC): A framework for perception and action planning, *Behavioral and Brain Sciences* 24 (2001) 849-878.
- G.R. Hunt, Manufacture and use of hook-tools by New Caledonian crows, *Nature* 379 (1996) 249-251.
- A. Iriki, O. Sakura, The neuroscience of primate intellectual evolution: Natural selection and passive and intentional niche construction, *Philosophical Transactions of the Royal Society B* 363(1500) (2008) 2229-2241.
- D. Kraft, E. Bacsieski, M. Popovic, N. Krüger, N. Pugeault, D. Kragic, S. Kalkan, F. Wörgötter, Birth of the object: Detection of objectness and extraction of object shape through object action complexes, Special issue on "Cognitive Humanoid Robots" in the *International Journal of Humanoid Robotics*, 2007 (in press).

- N. Krüger, M. Lappe, F. Wörgötter, Biologically motivated multi-modal processing of visual primitives, *The Interdisciplinary Journal of Artificial Intelligence and the Simulation of Behaviour* 1 (5) (2004) 417-428.
- G. Lakoff, M. Johnson, *Philosophy in the Flesh: The Embodied Mind and its Challenge to Western Thought*, Basic Books, New York, 1999.
- M. Lungarella, G. Metta, R. Pfeifer, G. Sandini, Developmental robotics: A survey, *Connection Science* 15 (4) (2003) 151-190.
- H. Maturana, F.J. Varela, *Autopoiesis and Cognition: The Realization of the Living*, Reidel, Dordrecht, 1980.
- J. McCarthy, P.J. Hayes, Some philosophical problems from the standpoint of artificial intelligence, *Machine Intelligence* 7 (1969) 195-204.
- D.J. McFarland, *Problems of Animal Behaviour*, Harlow: Longman Scientific & Technical, 1989.
- N. Mendes, D. Hanus, J. Call, Raising the level: Orangutans use water as a tool, *Biology Letters*, doi:10.1098/rsbl.2007.0198, 2007.
- S. Nolfi, D. Floreano, *Evolutionary Robotics, the Biology, Intelligence, and Technology of Self-Organizing Machines*, MIT Press/Bradford Books, 2000.
- S. Obayashi, M. Tanaka, A. Iriki, Subjective image of invisible hand coded by monkey intraparietal neurons, *NeuroReport* 11 (2000) 3499-3505.
- R. Pfeifer, C. Scheier, *Understanding Intelligence*, Cambridge, MA, MIT Press, 1999.
- J. Piaget, *The Child's Conception of Physical Causality*, Routledge and Kegan Paul, 1930.
- B. Porr, F. Wörgötter, Isotropic sequence order learning, *Neural Computation* 15 (2003) 831-864.
- B. Porr, F. Wörgötter, Inside embodiment – What means embodiment for radical constructivists? *Kybernetes* 34 (2005) 105-117.
- R.F. Port, T. van Gelder, *Mind as Motion: Explorations in the Dynamics of Cognition*, Cambridge, MA, MIT Press, 1995.
- D.J. Povinelli, *Folk Physics for Apes*, Oxford University Press, Oxford, 2000.
- N. Pugeault, Early cognitive vision: Feedback mechanisms for the disambiguation of early visual representation, Ph.D. Thesis, University of Göttingen, <http://webdoc.sub.gwdg.de/diss/2008/pugeault/>, 2008.
- A. Riegler, When is a cognitive system embodied? *Cognitive Systems Research* 3 (2002) 339-348.
- J.R. Searle, Minds, brains and programs, *Behavioral and Brain Sciences* 3 (1980) 417-457.
- Special Issue: *Trends in Cognitive Sciences* 10 (7) (2006).

- L. Steels, R. Brooks, *The Artificial Life Route to Artificial Intelligence: Building Embodied, Situated Agents*, Hillsdale, NJ, Erlbaum, 1995.
- R.S. Sutton, A.G. Barto, *Reinforcement Learning: An Introduction*, Bradford Books, MIT Press, Cambridge, MA, 2002 edition, 1998.
- J.B. Tenenbaum, T.L. Griffiths, C. Kemp, Theory-based Bayesian models of inductive learning and reasoning, *Trends in Cognitive Sciences* 10 (7) (2006) 309-317.
- E. Thelen, L.B. Smith, *A Dynamic Systems Approach to the Development of Cognition and Action*, Cambridge, MA, MIT Press, 1994.
- E.L. Thorndike, *Animal Intelligence*, New York, Macmillan, 1911.
- S. Thrun, W. Burgard, D. Fox, *Probabilistic Robotics*, MIT Press, 2005.
- S. Todes, *Body and World*, MIT Press, Cambridge, MA, 2001.
- F. Varela, E. Thompson, E. Rosch, *The Embodied Mind*, MIT Press, Cambridge, MA, 1991.
- H. von Foerster, On self-organizing systems and their environments, in: M. Yovits, S. Cameron (Eds.), *Self-Organizing Systems*, Pergamon Press, London, 1960, pp. 31-50.
- E. von Glasersfeld, Learning and adaptation in constructivism, in: L. Smith (Eds.), *Critical Readings on Piaget*, Routledge, London and New York, 1996, pp. 22-27.
- W.G. Walter, *The Living Brain*, G. Duckworth, London, 1953.
- A.S. Weir, J. Chappell, A. Kacelnik, Shaping of Hooks in New Caledonian Crows, *Science* 297 (2002) 981.
- J. Weng, J. McClelland, A. Pentland, O. Sporns, I. Stockman, M. Sur, E. Thelen, Artificial intelligence: Autonomous mental development by robots and animals, *Science* 291 (5504) (2001) 599-600.
- M. Wilson, Six views of embodied cognition, *Psychonomic Bulletin & Review* 9 (4) (2002) 625-636.

9 Appendix A: Body Extension based on Rigid Body Motion

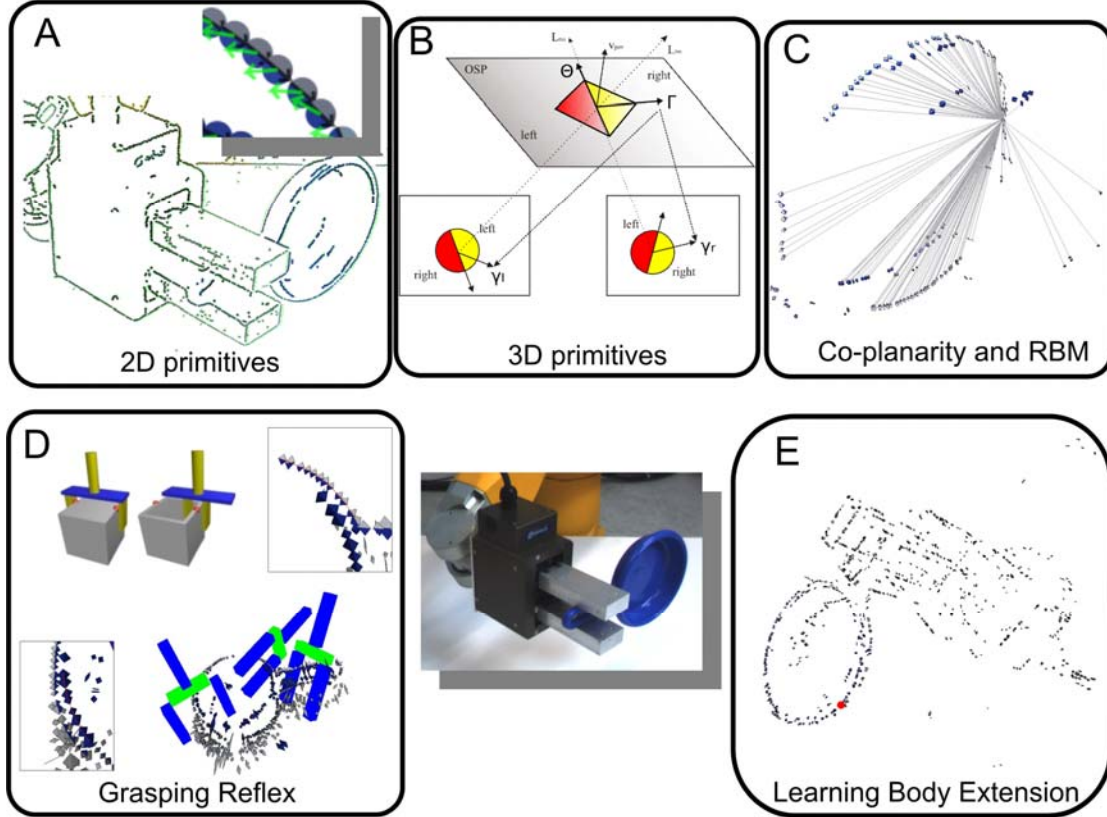


Fig. 6) Illustration of the scheme: A) 2D primitives are extracted from the image; B) stereo pairs of 2D primitives lead to the reconstruction of 3D primitives; C) coplanar pairs of 3D primitives are detected; D) grasp hypotheses are inferred from co-planar pairs; and E) the body extension takes place.

Local multi-modal Scene descriptors

The body extension algorithm makes use of visual symbolic descriptors called Multi-Modal Primitives. These primitives are extracted from images, and embedded in a hierarchical representation discussed in Krüger et al., (2004) and Pugeault (2008). 2D primitives (Fig. 6A) are defined as:

$$\pi = (x, \theta, \omega, (c_l, c_m, c_r)),$$

where x is the primitive's position in the image; θ is the 2D orientation; ω the local phase; and the colour is coded as three vectors (c_l, c_m, c_r) , corresponding to the primitive's left (c_l), middle (c_m) and right sides (c_r). Figure 6A shows the primitives extracted from an example scene and 6B the abstract representation of 2D (bottom) and 3D (top) primitives.

These primitives can be matched between two images taken from a pair of stereo cameras, and stereo pairs of 2D primitives allow reconstructing so-called 3D primitives that have the following formulation:

$$\Pi = (X, \Theta, \Omega, (C_l, C_m, C_r))$$

where X is the 3D position; Θ is the 3D orientation; and, Ω is the phase. Appearance based information is coded by generalising local phase and colour of the two corresponding 2D primitives. The reconstruction of a 3D primitive from two corresponding 2D primitives is illustrated in Fig. 6B. See Krüger, et al. (2004) and Pugeault (2008) for more information about 2 and 3D primitives and the extraction of the encoded modalities.

Relations can be defined on 3D primitives, of which two are relevant in the context of this paper. In the first place, a co-planarity relation between 3D primitives (see Fig. 6C) indicates to which degree these primitives are in the same plane based on the fact that the two 3D position and 3D orientations are an over-determined system. The other relation, called Rigid Body Motion (RBM), describes the transformation undergone by 3D primitives describing a moving rigid object, e.g., when the robot has grasped the object firmly and is manipulating it. For details, see Kraft et al. (2007) and Pugeault (2008).

Grasping Reflex based on Co-Planarity

Co-planar pairs of 3D primitives are detected using the co-planarity relation. Different predefined grasp types exist for different coplanarity relations. Two are shown at the top left of Fig. 6D. Then different grasping hypotheses can be associated to each actually existing co-planar pair in the image based on a best-match measure. Three of them are shown in Fig. 6D, bottom. In Kraft et al. (2007) we could show that, although these hypotheses do not always produce a successful grasp, many objects could be grasped by allowing multiple trials with different hypotheses of which eventually one will lead to a successful grasping action. The grasping success can be computed using haptic information as well as by measuring the gap between the grasper's two fingers after the grasp and after lifting the object.

Body knowledge based on the Rigid Body Motion Principle

The motion performed by an industrial robot (such as the Stäubli RX 60 that has been used in the experiments) is predicted with a high degree of precision by the motor control. Hence, the change undergone by 3D primitives, based on the RBM relation mentioned above, can also be predicted accurately. Then the body can be defined as the set of all 3D primitives that move according to this predicted motion. If the gripper fails to grasp anything, these primitives describe the gripper alone; if, on the other hand, the grasp has led to a stable physical control over an object, these primitives describe the gripper and the manipulated object (see Fig. 6E).

Body extension

The mechanisms described above provide the robot with a 3D model of its own body, including any extension obtained by the acquisition of stable physical control over an object (indeed the gripper itself can be seen as such an extension). Hence the ability to make contact with objects can be computed based on that knowledge and thus any grasped object extends the robot's workspace for (at least) simple actions such as pushing, as shown in Fig. 4.

10 Appendix B: Discovery by Doing

This appendix contains a more detailed description of the procedures used for the experiments shown in Fig.5 and of the example application implemented in a Stäubli robot (see also Agostini et al., 2008a,b).




Fig. 7) Example application using a Stäubli Arm.

Setup

Figure 7 shows our simple real world application implemented in a Stäubli arm (A) together with the used detector grid (B) and one example of a set of possible actions (C). The application consists of an environment with 14 cells configured in a 5 by 3 grid world where each cell in the grid represents a detector d_i for perception of grid occupancy. Bowls are placed in the grid and have black markers to ease the perception problem. Detector values are “black” when a bowl is centred in a cell and “white” otherwise. These values are derived from the originally existing average grey level in a given grid cell after thresholding. The goal is to free a certain area of the table to allow for moving one bowl (marked here with yellow rim, A) to the far end. Bowls are larger than grid cells, which needs also to be taken into account. To this purpose the arm can move any bowl, when possible, from its current cell to one of the neighbouring cells. Only vertical or horizontal movements of between one and four steps length are permitted (C), but diagonal ones could also be included if desired. Action are encoded by grid cell number (e.g. 7), direction of movement (e.g. U for up) and step length (e.g. 2 for 2 steps), leading to an encoding of, in this example, 7U2 (see Fig. 7C).

Formalism: Notation

We assume that the agent has a set of N sensors that measure some characteristics of the environment. The value of the sensor i is called an observation v_i . Each of these sensors is internally represented by the agent as a detector d_i (grid cell) that could take different values d_{ij} called *conditions* depending on the observation v_i . The function that maps observations to conditions is called the *perception function (PF)*, which basically transforms the raw information of the sensor into some “symbolic” representation manageable by the robot (Agostini et al., 2008b). Here this is a simple thresholding procedure on the measured grey levels.

An internal state s of the agent is constituted by the momentarily existing set of values of every detector, hence, $s = \{d_{1j}, d_{2k}, \dots, d_{Nl}\}$, where j, k, l are used to abbreviate possible detector values. In our case we have only two possible values for each detector (black,

white). Finally, we assume that at each moment the agent is capable of performing a well defined action a_i taken to the set of all possible actions A .

Formalism: Cause-effect Description – OACs

We represent a cause-effect description, in our words: an OAC, using a tuple that consists of a set of *pre-conditions*, an *action* a_i and a set of *post-conditions*. Pre-conditions are captured by the state s before the action rendering the description O of an object, post-conditions by the state s' after the action, rendering the description O' of the object after some of its attributes have changed. On top of this every OAC contains models of expected outcome measured as the Average Change $\langle \Delta O \rangle$ together with the assessment of the Repeatability $\sigma_{\Delta O}$ for the respective OAC and a Success Threshold Φ_S . Figure 5 shows multiple examples of different OACs where final detector values are shown after thresholding. Dashed cells (see Fig. 5) mean “don’t care” if there is either a black or a white.

Formalism: Expected Outcome

The expected outcome is coded as a set of confidence intervals I_i , one for each detector d_i involved in the *post-conditions* part of the OAC. The confidence intervals are built using a sampled mean $\langle \Delta O \rangle_i$ and standard deviation $\sigma_{\Delta O_i}$ of the observed condition values, the number of observations n_i that fed the statistics, and the sampled distribution t in the following way,

$$I_i = [m_i - t_{\alpha/2}^{n_i-1} \sigma_{\Delta O_i}, m_i + t_{\alpha/2}^{n_i-1} \sigma_{\Delta O_i}] = [I_i^{min}, I_i^{max}]$$

where $t_{\alpha/2}^{n_i-1}$ is the value of the t distribution with $p=\alpha/2$ and $n-1$ degrees of freedom.

The statistics are fed by the observations obtained after the OAC execution and if no failure occurs. If the number of samples n_i is below a predefined threshold n_c the teacher supervises the statistic updating as the agent doesn’t have enough confidence to evaluate by itself a failure.

Failures - Surprises

Several types of failures can occur in this system. These failures are fundamental in the teaching-learning process as they are the driving force for the agent “to ask for help”. This is evident as soon as a failures leads to the situation that a certain expected outcome did not take place, which is a typical “surprising situation”.

Specifically:

- 1) Failure happens, if any of the conditions obtained after an OAC execution lie outside their corresponding confidence intervals, i.e. if $P = \text{abs}(\Delta O - \langle \Delta O \rangle) > \Phi_S$, where $\Phi_S = (I_{max} - I_{min})/2$. As mentioned before, outcome evaluation only takes place after gathering enough statistics. In the special case of the application of Fig. 5 detector values are binary separated into black and white, hence convergence of the mean $\langle \Delta O \rangle$ will also have to be to black or white, and the standard deviations $\sigma_{\Delta O}$ are supposed to converge to 0. A failure will occur simply when the agent expects a white and perceives a black or vice versa. Surprises are solved by teacher explanations about what to “pay attention to” in order to obtain the expected outcome. In the approach explanations have the form of “detector A need to have value B” in order to obtain the desired outcome. With this explanation the robot corrects the OAC representation as depicted in the pseudo-code above.

- 2) The robot perceives the world incorrectly. Due to light changes and intrinsic uncertainties of the real world the robot could detect a bowl (black) in an empty cell or vice versa. As the teacher supervises robot perceptions with respect to state representations she will explain which is the proper attribute and the agent performs an updating of the perception function using the instructed attribute and the value obtained from the corresponding sensor.
- 3) Finally, a failure could also occur if there are no OACs that could produce the needed changes to reach the goal. In this case the teacher instructs about the action to be performed and the system generates a first approximation of the OAC with the observed changes.

Method

A complete explanation of the method developed for “discovery by doing” using OACs is beyond the scope of this article. See Agostini et al. (2008a,b) for a detailed description. Nevertheless, we present the basics of the method in the following pseudo-code, where we use the §§ symbol to denote a jump to a subroutine. Clearly this code fragment is not complete as many required checks, etc. are not been included for brevity.

Some explanations: *Goal* is the final state that the agent is supposed to reach by performing a sequence of OACs (see Fig. 5). *Rule* is a sequence of OACs by which the agent can link the perceptual prior perceived at the start to the final perceptual posterior at the goal. *Task* can be understood as the index of a certain OAC in a given rule. Hence when a certain task arises then one specific OAC has to be executed.

```

INIT system
Define GOAL
WHILE GOAL is not reached
{
  Perceive and store prior state Sprior using the "Perception Function" (2)
  Select a rule that connects Sprior to GOAL
  IF rule found
  {
    Extract all Tasks from the Rule to reach the GOAL
    FOR each Task
      §§Execute OAC(Sprior,Task)§§
  }
  ELSE (no rule)
  {
    Teacher instructs actions (1)
    FOR each action instructed
    {
      Perceive and store prior state Sprior (2)
      Execute action
      Perceive and store posterior state Spost (2)
      GENERATE new OAC using Sprior, Action and Spost
    }
    GENERATE RULES using new OACs
  } (end if rule found)
} (end while GOAL is not reached)

```

Execute OAC

```
Subroutine: Execute and update OAC(SPrior, Task)
  Execute action (deduced from SPrior, Task or instructed by teacher)
  Perceive posterior SPost
  Measure Change  $\Delta O$  (Posterior - Prior)
  Calculate  $U = \text{abs}(\Delta O - \langle \Delta O \rangle)$  (UNPREDICTABILITY)
  IF not the expected outcome, hence  $U > \Phi_S$  (FAILURE=SURPRISE)
  {
    Teacher explains bad/missing conditions (with indications like "detector A needs to have value B"
    to afford the expected outcome.)
    Correct OAC using teacher explanations (by replacing the value of detector A with
    attribute B in the OAC representation.)
  }
  ELSE (if not failure = success)
    Update OAC statistics: update  $\langle \Delta O \rangle$ ,  $\sigma_{\Delta O}$ , and  $\Phi_S$ 
Return
```

Actions of the Teacher

If necessary, the teacher interferes at three stages, see number in the pseudo-code: (1) she can instruct actions, (2) she can supervise whether a prior is correctly perceived or not, and (3) she can explain bad or missing conditions for an OAC.

These actions take place, if required, as described in the subsection on Failures – Surprises above.