A Standard Model of Mind:
AAAI Technical Report FS-17-05

The Distributed Adaptive Control Theory of the Mind and Brain as a Candidate Standard Model of the Human Mind

Paul F. M. J. Verschure

Synthetic Perceptive, Emotive and Cognitive Systems Laboratory; Catalan Institute of Advanced Studies (ICREA);
Catalan Institute for Bioengineering (IBEC); Center for Autonomous Systems and Neurorobotics,
Universitat Pompeu Fabra. Barcelona, Spain. specslab.com

Abstract

This article presents the Distributed Adaptive Control (DAC) theory of mind and brain as a candidate standard model of the human mind. DAC is defined against a reformulation of the criteria for unified theories of cognition advanced by Allen Newell, or the Unified Theories of Embodied Minds – Standard Model benchmark (UTEM-SM) that emphasizes real-world and real-time embodied action. DAC considers mind and brain as the function and implementation of a multi-layered control system and addresses the fundamental question of how the mind, as the product of embodied and situated brains, can obtain, retain and express valid knowledge of its world and transform this into policies for action. DAC provides an explanatory framework for biological minds and brains by satisfying well-defined constraints faced by theories of mind and brain and provides a route for the convergent validation of anatomy, physiology, and behavior in our explanation of biological minds. DAC is a well validated integration and synthesis framework for artificial minds and exemplifies the role of the synthetic method in understanding mind and brain. This article describes the core components of DAC, its performance on specific benchmarks derived from the engagement with the physical and the social world (or the H4W and the H5W problems) and lastly analyzes DAC's performance on the UTEM-SM benchmark and its relationship with contemporary developments in AI.

Introduction

Allen Newell proposed to address the "great psychological data puzzle" by postulating a single set of mechanisms for all cognitive behavior or Unified Theories of Cognition (UTC). He devised a list of criteria that any UTC had to satisfy: 1) Behave flexibly as a function of the environment; 2) Exhibit adaptive (rational, goal-oriented) behavior; 3) Operate in real-time; 4) Operate in rich, complex, detailed environments; 5) Use symbols and abstractions; 6) Use language; 7) Learn from the environment and from experience; 8) Acquire capabilities through development; 9) Operate autonomously, but within a social community; 10) Be self-aware and have a sense of self; 11) Be realizable as a neural system; 12) Be constructible by an embryological growth process; 13) Arise through evolution (Newell, 1994, p.19). We can take this list as a starting point to define the criteria a standard model of the human mind should satisfy (See also Anderson & Lebiere, 2003). Newell advanced his SOAR theory as a candidate UTC. Although Newell made an important step in the advancement of UTC's and achieving human level performance is again a major theme in contemporary AI, no generally accepted UTC is available today. I propose that this is due to the insufficient considerations given to embodiment, situatedness, computation and learning in the development of mind. This is based on developments in contemporary AI, cognitive science and advances in neuroscience in particular concerning the system level organization of the embodied brain. Hence, now about 25 years later we can recalibrate UTC by redefining as criteria for Unified Theories of *Embodied* Minds – Standard Model (UTEM-SM) where candidate UTEM-SMs have to answer both functional and structural constraints:

1: Functional constraints (Psychology of mind):

Level 1: Display autonomous adaptive and flexible real-time goal-oriented behavior in complex physical environments (Newell test: 1, 2, 3, 4, 7, 10-sense of self);

Level 2: Display autonomous adaptive and flexible real-time goal-oriented behavior in complex real-world social environments including the use of symbols and language (Newell test: Level 1 + 5, 6, 9, 10-self-aware);

2: Structural constraints (Biology of embodied brain):

Biological validity: be plausibly the product of biological evolution and be demonstrably constructible through neuro- and morphogenesis (Newell test: 11-13)

Physical realizability: perform in real-time, in the real-world using resources (e.g. energy, computation) comparable to biological systems.

The UTEM-SM benchmark refines and elaborates Newell's original UTC benchmark by defining levels of performance (physical vs social), psychological and biological validity and by including behavioral and epistemic autonomy together with real-world and real-time performance as specific constraints. This implies that UTEM-SM argues that UTCs should be *both* scientific theories of natural minds and brains as well as models for synthesizing artificial ones. In this way the misleading effects of mimicry that the Turing test faces is avoided while biological inspiration is rejected as a method.

Distributed Adaptive Control (DAC) as a candidate standard model of the human mind

Distributed Adaptive Control (DAC) is a theory of mind and brain introduced in 1991 against the background of developments in the fields of artificial intelligence and cognitive science and has been generalized towards robotics, cognitive science, psychology, neuroscience, neurology and education. DAC departs from the fundamental paradox between rationalism and empiricism: does knowledge originate in reason as advanced by Plato or in the senses as proposed by his student Aristotle? At the heart of the contradiction between these schools of thought stands the obstacle of the putative inability of scaling sensory data to valid knowledge and reasoning as assumed by Descartes and argued by Hume. This dilemma has dominated the study of mind and brain in the 20th century, for instance in the criticism leveled by Chomsky against Skinner's theory of language, triggering the AI revolution of the 1950ies and by cognitivists such as Fodor against the connectionism of the 1980ies. There is still no consensus on how this fundamental problem of epistemology can be resolved. DAC advances a theory of this knowledge problem based on a notion of constructive empiricism where scientific theories strive towards being empirically adequate (Van Fraassen, 1980) through the convergent validation of linking to the multiple levels of description of biological and artificial minds and brains (Verschure, 2012). However, given that we can see the mind/brain as a knowledge organ, DAC has taken this fundamental epistemological challenge as its point of departure. In AI this problem is also known as the symbol grounding problem or how can an artificial intelligence conceived as a Turing machine assign meaning to its symbols? (Harnad, 1990). In a broader perspective, we can speak of the problem of priors or what is the minimal set of rules and representations to bootstrap a natural or artificial intelligence? (Verschure, 1998;2012). The driving intuition behind DAC is that knowledge is grounded in the interaction between embodied agents and the physical and social environments in which they are situated, constrained by mechanisms of learning and memory and bootstrapped from minimal priors that emerged during evolution. DAC was of immediate relevance to the nascent field of "New AI " (See for a review Verschure, 2012) and it was also one of the first neural models of a cognitive architecture ported to both simulated and physical robots (Mondada & Verschure, 1993).

DAC takes 'mind' as the collections of functional macroscopic properties of embodied brains that are directly or indirectly expressed in action. Mind is an amalgamation of processes supporting motivation, perception, cognition, attention, memory, learning, action and consciousness. The mind is situated in physical and social environments, and because of the tight coupling of body, brain, mind, and environment especially when taking into account memory, we can speak of a nexus (Verschure, 2012). 'Behaviour' is defined as autonomous changes in the position or shape (confirmation) of the body or soma of an agent. Once behavior serves internally-generated goals we can speak of 'action', i.e., it is intentional or conative. The 'brain' is a distributed, wired control system that exploits the spatial organization of connectivity combined with the temporal response properties of its units to achieve transformations from sensory states, derived from both the internal (body) and external (world) environment, into action. DAC follows 19th-century physiologists Claude Bernard and Ivan Pavlov in conceptualizing the mind/brain as a control system that generates action to maintain a multi-stable equilibrium between the body and the environment. DAC thus downgrades the importance of dominant metaphors such as information processing or computation. In other words, information is created and processed, and processes realized that might be describable in computational terms of symbol manipulation, but these are at best descriptions of processes and mechanism that serve and are predicated on the realization of current and the planning of future goal oriented action by an embodied and situated agent. The core variable the DAC mind/brain controller maintains in a dynamic equilibrium is the integrity of the organism in the face of the second law of thermodynamics, defined through the organism's needs that are continuously challenged by somatic and environmental change. A control system is usually considered as distinct from what it controls, i.e., the plant. However, DAC considers the body as an integral part of the control system itself (Verschure & Pfeifer. 1992; Verschure et al, 2003).

From the perspective of control, the question is what top-level functions and essential variables the mind/brain optimizes to generate the How of action? DAC proposes that these are just five: needs, drives and motivation or "Why"; states of the world such as objects or "What"; spatial structure of the task or "Where"; the temporal dynamics of the task and the agent or "When"; "Who" in case the agent deals with other agents. The function of the mind/brain is to solve the H4W optimisation problem when a single agent confronts its physical world and H5W in case the world also contains other agents such as predators, prey, and conspecifics (Verschure, 2016).

The structure of DAC

DAC comprises four coupled layers of control (Figure 1). *The Somatic Layer* (SL) of DAC designates the body and defines three fundamental sources of information: sensation driven by external and internal sources of stimulation, needs defined by the essential variables that assure survival, and actuation defined by the control of the skeletal-muscle system.

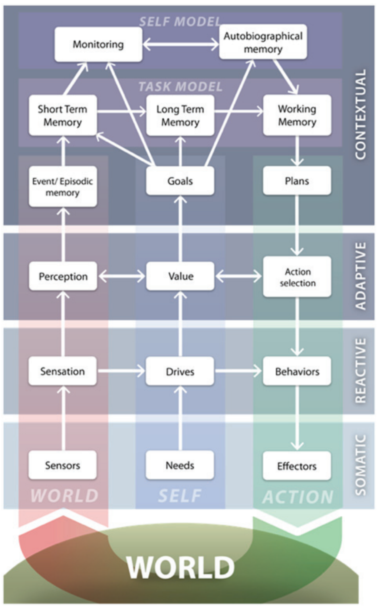

Figure 1. Abstract representation of the Distributed Adaptive Control (DAC) theory showing its main processes (boxes) and dominant information flows (arrows). DAC comprises four layers (Soma, Reactive, Adaptive and Contextual) and three functional columns: exosensing, the sensation and perception of the world (left,red) endosensing, detecting and signalling states derived from the physically instantiated self (middle, blue); and the interface between self and the world through action (right, green). The arrows show the primary flow of information mapping exo- and endosensing into action, defining a continuous loop of interaction via the world. The Soma designates the body with its sensors, organs, and actuators defining the needs, or Self Essential Functions (SEF), the organism must satisfy to survive. The Reactive Layer (RL) comprises dedicated Behaviour Systems (BS) each implementing predefined sensorimotor mappings serving the SEFs. A so called allostatic controller regulates action selection, task switching, and conflict resolution among all BSs by setting their internal homeostatic dynamics relative to overall system demands. The Adaptive Layer (AL) acquires a state space of the agent-environment interaction and shapes action. The learning dynamics of AL is constrained by the value functions derived from the allostatic control of the RL and minimises perceptual and behavioral prediction error building a model free action generation system. The contextual layer (CL) expands the time horizon in which the agent can operate, realizing model based policies, through the use of sequential short and long-term memory systems (STM and LTM respectively). STM acquires conjunctive sensorimotor representations that are generated by the AL as the agent acts in the world. STM sequences are retained as goal-oriented model based policies in LTM triggered by value signals driven by the RL and AL. The contribution of the LTM policies to goal oriented decision-making depends on four factors: goal states, perceptual evidence, memory chaining, valence (including the expected cost of reaching a given goal state). The content of working memory (WM) is defined by the memory dynamics that represents this four-factor decision-making model. The autobiographical memory system of CL allows the restructuring of memory around the unifying notion of Self which is of particular relevance for the interaction with the social world. See text for further explanation.

The Reactive Layer (RL) [8] comprises fast predefined sensorimotor loops that support direct behaviors underlying Self Essential Functions. These reflexes are coupled via need and drive systems creating sense-affect-act triades. In contrast to standard reactive models (e.g., Brooks, 1986), the distinguishing feature of the DAC RL is that it is part of a larger architecture and serves distinct behavioral and epistemic functions. The activation of a reflex carries essential information on the interaction between the agent and the world that is a key control signal for subsequent layers driving conflict resolution and epistemic needs, i.e., knowledge acquisition (See below). To avoid conflicts between internal states, such as avoidance and approach, a competitive relationship exists between the different internal states. A further distinguishing feature of this layer is that it is modeled in terms of an allostatic process that regulates the homeostatic behavior sub-systems serving SEFs. This is both closer to the dynamics of physiological systems and scalable as opposed to more phenomenological behavior found in behavior-based robotics (e.g.,(Arkin, 1998). The SEFs of RL are both oriented towards direct survival as well as epistemic functions such as exploration and novelty seeking.

The Adaptive Layer (AL) extends the predefined sensorimotor loops of the RL with acquired sensory and action states. Hence, it allows the agent to transcend from strictly predefined reflexes through learning. The AL is interfaced to the full sensorium of the agent, its internal needs and its effector systems receiving internal state information from the RL and in turn generating action. AL comprises adaptive mechanisms to deal with the fundamental unpredictability of both the internal and the external environment, e.g. the symbol grounding problem. Through learning, a state space of world states is acquired together with the shaping of action patterns and their association. The AL models the learning dynamics of classical conditioning advancing a prediction based Hebbian learning rule Verschure & Pfeifer, 1992) which has been phrased in a general formal framework: correlative subspace learning (CSL) (Duff & Verschure, 2010). The AL of DAC provides a solution to the problem of priors/symbol grounding described earlier because it acquires the state space of the world and the agent through its interaction with the environment. CSL captures the law of associative competition formulated by Rescorla and Wagner that emphasizes that learning depends on how unexpected a stimulus is given the internal state of the organism (Rescorla & Wagner, 1972). The CSL model is also consistent with adaptive filter methods going back to the Kalman filter and other derived approaches (Kalman, 1960). Indeed, DAC is an early example of the "predictive brain" perspective (Friston, 2010;Maffei, et al., 2014). The AL allows the agent to overcome the predefined behavioral repertoire of the RL and to successfully engage with unpredictable aspects of the world.

The Contextual Layer (CL) of DAC expands the spatio-temporal window of action by developing policies for goal-oriented action using systems for short-, long- and working-memory (Figure 2). These memory systems allow for the formation of sequential representations of states of the environment and actions generated by the agent. The atomic elements are formed by the state space of exo- and endo-sensing constructed by the AL or its sensorimotor contingencies. The acquisition and retention of these sequences are conditional on the goal achievement of the agent. These behavioural plans can be recalled through sensory matching and internal chaining among the elements of the retained memory sequences. The dynamic states that this process entails are DAC´s working memory system.

Experiments with the CL of DAC with simulated and physical robots revealed that a unique feedback loop exists between action and perception that stabilizes the interaction between the AL and CL through behavior itself called behavioral feedback (Verschure et al., 2003). Thus model based policies or behavioral plans acquired by the CL, through their mapping to action carve out an effective behavioral space or niche rendering the world more predictable and the perceptual reconstruction error smaller.

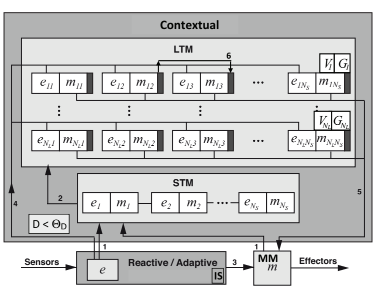

Figure 2: Contextual layer of DAC: (1) The predicted perceptual states or prototypes and the motor activity generated by the AL are acquired and stored as a segment in STM if the discrepancy between predicted (e) and encountered (x) sensory states, D, falls below a predefined threshold. D is the time averaged reconstruction error of the AL perceptual learning system: x - e; (2) If a goal state is reached, e.g., reward or punishment detected, STM is retained in LTM conserving its order, and STM is reset. Every sequence is labeled with the specific goal and internal state/valence it pertains to; (3) the motor population (MM) receives input per the rules of the AL; (4) If RL/AL motor activity is sub-threshold, the values of the current prototypes, e, are matched against those stored in LTM and action options are generated by optimizing goals, perceptual evidence, value and memory chaining. (5) MM receives the CL action policy derived motor response as a weighted sum over the active memory segments of LTM and performs action selection. (6) The segments that contributed to the executed action will prospectively bias other associated LTM segments. The Contextual Layer (CL) bootstraps the system further to deal with novel and a priori unknown states of the world and the agent in an extended spatio-temporal window creating behavioral plans or policies (Figure 2). The CL comprises systems for short and long-term and working memory (STM, LTM, and WM, respectively). These memory systems allow for the formation of sequential representations of states of the environment and actions generated by the AL or its acquired sensory-motor contingencies. The acquisition and retention of these sequences are conditional on the goal achievement of the agent and the absence of RL activation. CL policies are recalled through sensory matching and internal chaining among the elements of the retained memory sequences. The dynamic states that this process entails are the CL's working memory system.

A further question was whether the solutions that DAC found could be considered optimal. Lacking consensus benchmarks and comparable approaches a more formal approach was taken based on work by Massaro who proposes based on his studies of multi-modal speech perception that Bayesian integration is a universal principle of perception and cognition (Massaro, 1997). Indeed, DAC comprises constructs that are analogs of the central components of a Bayesian analysis: goals, actions, hypotheses, observations, experience, prior probabilities and a score function. By phrasing the robot foraging tasks performed with the DAC architecture in Bayesian terms, it was shown that the CL generates actions that are optimal in a Bayesian sense (Verschure & Althaus, 2003). In this perspective, the optimal action, *a*, is the one that optimises the expected gain $\langle g \rangle_a$:

$$\langle g \rangle_a = \sum_{s_n \in S} p(s_n|r) G_g(s_n, a)$$

where $p(s_n|r)$ is the posterior probability defined by Bayes' rule and $G_g(s_n, a)$ is a score function that defines the gain obtained from executing action *a* if s_n is true. *s* and *r* are the perceptual predictions generated and stored at the contextual and adaptive layers, respectively and G is determined by the labeling of the LTM sequences regarding the goal states they are associated with. These goal states are the top-level representations generated by the Self column of the architecture comprising needs, drives, value and goals and their associations with states of the world. The actions selected by the CL are Bayesian optimal with respect to G. Hence, DAC can be considered an autonomous embodied rational system that acquires its own knowledge through its interaction with the environment and subsequently uses it in a Bayes optimal fashion to reach its goals, a key step towards UTEM-SM.

Answering the UTEM-SM benchmark with DAC

DAC has been widely used to address the H4W challenge as in robot foraging tasks (Level 1). However, to satisfy the UTEM-SM benchmark, it must also generalize to the social world, or H5W, (Level 2) and satisfy structural constraints derived from the brain and functional ones from psychology. DAC has been successfully mapped to the specifics of the neuroscience of mind and brain along two types of models to realize convergent validation. On the one hand, a whole brain architecture approach was followed which facilitates the mapping to behavior and psychology, while components of the architecture and basic operating principles have been linked to the invertebrate and mammalian brain through anatomically and physiologically constrained models [22]. These two lines of the DAC project have been integrated into a first embodied whole brain model comprising detailed models of core brain structures including cerebellum, entorhinal cortex, hippocampus and Prefrontal/Premotor cortex, or DACX (Maffei et al, 2015). DACX validates the overall DAC model in the context of foraging including obstacle avoidance, hoarding, exploration and homing. In parallel, predictions of DAC derived hypotheses have been validated in the laboratories of experimental neuroscientists and core principles of DAC have been mapped to highly effective neurorehabilitation and education technologies (See for a review Verschure, 2012). This shows the commitment of DAC to convergent validation with the goal to obtain an empirically adequate description of biological minds and brains and its validation through real-world applications.

To generalize from H4W to H5W, the DAC architecture has been mapped to the control of anthropomorphic robots that engage in dyadic interactions with humans (Lallée et al., 2015). This step has included the augmentation of the functions of the architecture to include drives to socially engage, to seek knowledge of the world through interaction with humans, the ability to acquire models of other agents and "read their minds", to use social cues to establish and maintain interaction and to learn language. Underlying the successful deployment of DAC in H5W scenarios the system was augmented with an autobiographical memory system that allows the robot to anchor its experiences in the H5W ontology (Lallee & Verschure, 2015). This DACh, h for Humanoid, architecture has been shown to successfully solve the H5W challenge in restricted task domains with a single human, another key step towards UTEM-SM. The current challenge is to scale the current H5W capabilities up to to make anything a task, the requirement for Robot Artificial General Intelligence and the DAC theory predicts that this will require a form of robot consciousness (Verschure, 2016).

Currently, artificial intelligence is in the grips of a third wave of neural network modeling after the early models of Rashevsky and Rosenblatt and their reprise in the connectionism of the late 1980ies. Two approaches stand out that again fall at opposite sides of the rationalism-empiricism divide. First, due to advances in computing technology and the availability of massive data sets, learning in multi-layered neural networks has made significant progress into complex task domains (Lecun, Bengio, & Hinton, 2015; Schmidhuber, 2015). Hence, deep learning assumes minimal priors but requires extremely large data sets for training. These data sets, however, are still labeled by humans to drive supervised learning. Subsequently, deep learning has been combined with reinforcement based learning with the goal to reach human or super-human performance in benchmark tasks such as Atari video games (Mnih et al., 2015) and complex games such as Go (Silver et al., 2016). These approaches still require large quantities of trials to train the network that is much higher than what humans need. In response, others have proposed that one-shot learning can be achieved when pre-existing core knowledge is provided in the form of physics and psychology simulation engines (typically Bayesian causal models (Lake et al., 2015). These approaches do express core tenets of the DAC theory such as the need for learning dependent acquisition of the state space, its compression and the combined realization of this state space with the development of action policies. However, they are also still problematic when applied to real-world systems. Deep learning approaches require large data sets made up of uncorrelated adjacent samples as well as significant computational resources, making them hardly compatible with the real-time constraints faced by real-world minds. Bayesian causal models, in turn, require extensive prior knowledge dependent on the physical and social environment, which still must be accounted for, is hardly generalizable across different applications and will remain brittle in the face of real-world variability. Both approaches are critically dependent on human labeling of data and pre-specification of prior knowledge and procedures and are thus failing on the problem of priors and symbol grounding among others. DAC advances and has implemented an alternative approach considering that embodiment and situatedness provide a grounded source of priors for embodied minds based on the following principles: First, embodiment strongly constrains the viable set of sensorimotor states of the agent through the specific physical coupling of the body and the environment. These constraints are specific to the particular morphology of an embodied mind, its internal needs (e.g. in term of energy consumption and damage minimization) and the environment in which it operates, precluding a full explicit pre-definition as assumed in terms of full pre-labeling of data and an intuitive physics/psychology engine. Second, situatedness places the embodied mind in the social environment governed by specific norms and conventions and provides crucial cues about well-adapted behavioral policies [24]. Here again, this is specific to the context in which embodied minds operations and cannot be fully predefined as assumed by an intuitive psychology engine. Also, DAC has taken real-world benchmarks of greater ecological validity such as foraging, language learning and real-world maze solving than computer and board games. These more contemporary approaches are essentially providing alternative views on sub-systems accounted for in the DAC theory (Moulin-Frier et al. , 2017) and consequently are still far removed from achieving the UTEM-SM-SM challenge. In contrast, DAC has already made significant inroads towards this benchmark by addressing both levels of functional constraints (H4W and H5W) and by having established strong links with the neuronal principles and psychological processes of perception, cognition, and action. However, to quote Allan Turing: "We can only see a short distance ahead, but we can see plenty there that needs to be done."

Acknowledgements

The DAC project has received support from a number of projects supported by the European Commission in particular SF (217418), EFAA (270490), GoalLeader (270108), eSMCs (270212), socSMCs (641321), WYSIWYD (612139), the European Research Council project cDAC (ERC 341196).

References

Anderson, J. R., & Lebiere, C. 2003. The Newell Test for a theory of cognition. *The Behavioral and Brain Sciences*, *26*(5):587-601.

Arkin, R. C. 1998. *Behavior-based robotics*. MIT Press.

Brooks, R. A. 1986. A robust layered control system for a mobile Robot. *IEEE: J. on Rob. and Aut.*, *ra.2*:14–23.

Duff, A., & Verschure, P.F. 2010. Unifying Perceptual and Behavioral Learning with a Correlative Subspace Learning Rule. *Neurocomputing*, *73*:1818–1830.

Friston, K. 2010. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, *11*(2):127–138.

Harnad, S. 1990. The Symbol Grounding Problem. *Physica D*, *42*:335–346.

Kalman, R. E. 1960. A new approach to linear filtering and prediction problems. *Tr. ASME J. of Basic Eng*, *82*:35–45.

Lake, B. M., Salakhutdinov, R., & Tenenbaum, J. B. 2015. Human-level concept learning through probabilistic program induction. *Science*, *350*(6266):1332–1338.

Lallee, S., & Verschure, P.F. 2015. How? Why? What? Where? When? Who? Grounding Ontology in the Actions of a Situated Social Agent. *Robotics*, *4*(2):169–193.

Lallée, S., Vouloutsi, V., Blancas, M., Grechuta, K., Puigbo, J., Sarda, M., & Verschure, P. F. M. . 2015. Towards the synthetic self: making others perceive me as an other. *Paladyn Journal of Behavioral Robotics*, *6*(1):136-64

Lecun, Y., Bengio, Y., & Hinton, G. 2015. Deep learning. Nature, 521:436-44

Maffei, G., Sanchez-Fibla, M., Herreros, I., & Verschure, P.F. 2014. *The role of a cerebellum-driven perceptual prediction within a robotic postural task. From Animals to Animats 13:76-87*

Maffei, G., Santos-Pata, D., Marcos, E., Sánchez-Fibla, M., & Verschure, P. F. 2015. An embodied biologically constrained model of foraging DAC-X. *Neural Networks*. 72:88-108

Massaro, D. W. 1997. *Perceiving Talking Faces : From Speech Perception to a Behavioral Principle*. MIT Press.

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., … Hassabis, D. 2015. Human-level control through deep reinforcement learning. *Nature*, *518*(7540):529–533.

Mondada, F., & Verschure, P.F. 1993. Modeling system-environment interaction. In *Proceedings of the Second European Conference on Artificial Life, Brussels* (pp. 808–817). MIT press.

Moulin-Frier, C., Puigbò, J.-Y., Arsiwalla, X. D., Sanchez-Fibla, M., & Verschure, P. 2017. Embodied Artificial Intelligence through Distributed Adaptive Control: An Integrated Framework. ICDL-EPIROB 17, In Press.

Newell, A. 1994. *Unified theories of cognition*. Harvard University Press.

Rescorla, R. A., & Wagner, A. R. 1972. A theory of Pavlovian conditioning. In *Classical Conditioning II Current Research and Theory. Eds Black AH, Prokasy WF. 64-99. Appleton Century Crofts*.

Schmidhuber, J. 2015. Deep learning in neural networks: An overview. *Neural Networks*, *61*:85–117.

Silver, D., … Hassabis, D. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature*, *529*(7587):484–89.

Van Fraassen, B. C. 1980. *The scientific image*. *Clarendon library of logic and philosophy* (Vol. 33).

Verschure, P. 2012. Distributed adaptive control: a theory of the mind, brain, body nexus. *Biologically Inspired Cognitive Architectures*. 1:55-72

Verschure, P.F. 1998. Synthetic Epistemology. In *IEEE World Congress on Computational Intelligence* 98:147–153.

Verschure, P.F. 2016. Synthetic consciousness: the distributed adaptive control perspective. *Philosophical Transactions of the Royal Society London. Series B*, *371*(1701):263–275.

Verschure, P.F., & Althaus, P. 2003. A real-world rational agent: unifying old and new AI. *Cognitive Science*, *27*(4):561–590.

Verschure, P.F. & Pfeifer, R. 1992. Categorization, representations, and the dynamics of system-environment interaction. *From Animals to Animat 2*:210–217. MIT press.

Verschure, P. F., Voegtlin, T., & Douglas, R. J. 2003. Environmentally mediated synergy between perception and behaviour in mobile robots. *Nature*, *425*:620–24.