

Computational Constructivist Model as an Anticipatory Learning Mechanism for Coupled Agent–Environment Systems

Filipo Studzinski Perotto • Constructivist AI Research Group, France • filipo.perotto/at/ufrgs.br

> Context • The advent of a general artificial intelligence mechanism that learns like humans do would represent the realization of an old and major dream of science. It could be achieved by an artifact able to develop its own cognitive structures following constructivist principles. However, there is a large distance between the descriptions of the intelligence made by constructivist theories and the mechanisms that currently exist. **> Problem** • The constructivist conception of intelligence is very powerful for explaining how cognitive development takes place. However, until now, no computational model has successfully demonstrated the underlying mechanisms necessary to realize it. In other words, the artificial intelligence (AI) community has not been able to give rise to a system that convincingly implements the principles of intelligence as postulated by constructivism, and that is also capable of dealing with complex environments. **> Results** • This paper presents the constructivist anticipatory learning mechanism (CALM), an agent learning mechanism based on the constructivist approach of AI. It is designed to deal dynamically and interactively with environments that are at the same time partially deterministic and partially observable. CALM can model the regularities experienced in the interaction with the environment, on the sensorimotor level as well, as by constructing abstract or high-level representational concepts. The created model provides the knowledge necessary to generate the agent behavior. The paper also presents the coupled agent environment system (CAES) meta-architecture, which defines a conception of an autonomous agent, situated in the environment, embodied and intrinsically motivated. **> Implications** • The paper can be seen as a step towards a computational implementation of constructivist principles, on the one hand suggesting a further perspective of this refreshing movement on the AI field (which is still too steeped in a behaviorist influence and dominated by probabilistic models and narrow applied approaches), and on the other hand bringing some abstract descriptions of the cognitive process into a more concrete dimension, in the form of algorithms. **> Constructivist content** • The connection of this paper with constructivism is the proposal of a computational and formally described mechanism that implements important aspects of the subjective process of knowledge construction based on key ideas proposed by constructivist theories. **> Key words** • **Factored partially observable Markov decision process (FPOMDP), computational constructivist learning mechanisms, anticipatory learning, model-based learning.**

Introduction

« 1 » The constructivist approach to artificial intelligence can be defined as the set of works on this science directly or indirectly inspired by ideas coming from the constructivist conception of intelligence. This conception was essentially defined by Jean Piaget (1954) and gave rise to an important school of thought that influenced many scientific fields from the second half of the twentieth century onward. The first important AI system based on constructivist concepts appeared much later, presented by

column A

Gary Drescher (1991), but even if his model had some theoretical impact on the field of AI, it could never be used to solve significant applied problems. Since then, year after year, new papers have been published that attempt to point out a way to implement such a strong mechanism (Guerin 2011). However, the constructivist approach has never thrilled most researchers in the AI community, staying in that uncomfortable position between the promise of true intelligence and the lack of impressive results.

« 2 » In this article, we present the *constructivist anticipatory learning mechanism*

column B

(CALM), an agent learning mechanism based on the constructivist approach of AI. CALM is designed to discover regularities in partially deterministic environments: it identifies the deterministic transformations present in non-deterministic situations. The mechanism operates incrementally: the agent learns at the same time as it needs to interact with the environment. CALM can also deal with partially observable environments: it is able to infer the existence of hidden or abstract properties, integrating them in its anticipatory cycle. The constructed anticipatory model can

column C

column A

1 be used by the agent to optimize its action
2 policy, improving its performance on its
3 own activities and adapting its behavior to
4 the experienced reality according to self-
5 determined goals or internal motivations.

6 « 3 » The elementary piece of knowl-
7 edge used by the mechanism is the *schema*,
8 an anticipatory structure that can be de-
9 scribed in the form of a conjunctive im-
10 plication: $context \wedge action \rightarrow expectation$. It
11 represents the prediction of experiencing
12 some transformation when a given action
13 is carried out in a given context.

14 « 4 » CALM is the cognitive engine
15 embedded in an artificial agent. In order to
16 complete the description of our construc-
17 tivist computational model, in this article
18 we present the *coupled agent environment*
19 *system* (CAES), a meta-architecture that
20 defines the agent as an autonomous entity,
21 situated in the environment, embodied and
22 intrinsically motivated.

23 « 5 » Our work aims to constitute one
24 more brick in the effort to bridge the gap
25 between the insightful but too abstract
26 descriptions of intelligence made by con-
27 structivist theories and robust artificial in-
28 telligence mechanisms able to implement
29 them.

30 Sensorimotor to symbolic

31
32
33
34 « 6 » The gradual development of a
35 symbolic intelligence over a sensorimo-
36 tor intelligence is an essential aspect of
37 explaining how human beings can render
38 intelligible their experiences, giving some
39 sense to the world, and learning to interact
40 with it (Piaget 1954). The challenge is the
41 same for a situated artificial agent (such as
42 a robot), who needs to learn incrementally
43 the regularities observed throughout its in-
44 teraction with the environment where it is
45 inserted.

46 « 7 » The experienced reality is some-
47 thing subjective and should not be con-
48 fused with an external, objective, ontologi-
49 cal universe, which is assumed to be on the
50 other end of the interaction interface. The
51 world as it is cannot be apprehended out-
52 side the domain of experience; whatever
53 may lie beyond sensorial perception is in-
54 accessible (Glaserfeld 1974, 1979). Any
55 representation of the outside reality will

column A

column B

be a model necessarily based on regulari-
ties extracted from subjective sequences of
observations and actions, and not from the
structure of that reality, which remains un-
known in its essence.

« 8 » Moreover, in complex environ-
ments, special “macroscopic” properties
emerge from the functional interactions
of “microscopic” elements, and such emer-
gent characteristics are not defined in any
of the sub-parts that generate them (Gold-
stein 1999). The salient phenomena in this
kind of environment tend to be related to
high-level objects and processes (Thornton
2003). In this case, if we suppose the exist-
ence of a complex universe out there, it is
plainly inadequate to represent the experi-
ence only in terms of primitive sensorimo-
tor elements (Drescher 1991).

« 9 » Considering these conditions, an
intelligent agent (human or artificial) must
have the capacity to overcome the limits of
pure sensorial perceptions, organizing the
universe in terms of more abstract concep-
ts. The agent needs to be able to detect
high-level regularities in the dynamics of
the environment, but this is not possible
if the agent is stuck in a rigid *representa-
tional vocabulary*.¹ In a constructivist ap-
proach, cognitive development must be a
process of gradual complexification of the
intelligence, where more abstract struc-
tures (symbolic) are built from simpler
sensorimotor interactions, in a way that
harmonizes the lived experiences with the
constructed model.

« 10 » From the flat, unstructured, con-
tinuous flow of perceptions resulting from
the situation of the agent in a complex uni-
verse, intelligence needs to build some or-
ganization. While the constructed internal
knowledge might reflect an external reality
to some degree, from the agent’s perspec-
tive this remains undecidable. Importantly,
though, intelligence progressively organ-
izes knowledge in increasingly abstract
structures, enriching the agent’s under-
standing of its own experiences.

1| The agent’s representational vocabulary
is the set of elements it can manipulate to create
knowledge.

column B

column C

1 Situativity, embodiment

2 and intrinsic motivation

3
4 « 11 » A given universe (natural or
5 computationally simulated) is a whole sys-
6 tem that can be analytically separated into
7 two different entities: an *agent* and an *envi-
8 ronment*. These two entities can be defined
9 as mutually dependent *dynamical systems*,²
10 partially open to each other, and continual-
11 ly deforming their trajectories (Beer 1995,
12 2004; Barandiaran & Moreno 2006; Ashby
13 1952).

14 « 12 » A situated agent (Wilson & Clark
15 2008) is an entity embedded in an environ-
16 ment. Due to the fact that the agent is only
17 one among many forces that generate the
18 environment dynamics, it is only partially
19 capable of transforming the environment
20 by its actions. In the same way, due to the
21 fact that the agent’s sensorial perception is
22 limited in some manner, the environment
23 becomes only partially observable and the
24 agent can find itself unable to distinguish
25 between differing states of the world (Such-
26 man 1987). The same situation can be per-
27 ceived in different forms, and different situ-
28 ations can have a similar appearance. This
29 ambiguity in the perception of states, also
30 referred to as *perceptual aliasing*, has seri-
31 ous effects on the ability of most learning
32 algorithms to construct consistent knowl-
33 edge and stable policies (Crook & Hayes
34 2003).

35 « 13 » The agent is embodied (Ander-
36 son 2003; Ziemke 2003): it presents inter-
37 nal states and metabolisms, elements that
38 belong neither to the mind nor to the en-
39 vironment. This characteristic allows the
40 agent to have intrinsic motivations: evalu-
41 ative signals related to the internal state of
42 the agent, and not to external environmen-
43 tal states to be reached.

2| A dynamical system consists of an ab-
stract state space evolving over time according
to a rule that specifies the immediate future state
given the current state.

column C

column A
Coupled agent–environment system

1
 2
 3
 4 « 14 » CAES is a meta-architecture
 5 proposed in this article to define a coupled
 6 agent–environment system, respecting the
 7 notions described in the precedent section.
 8 The *universe* (U) is represented as a global
 9 system $U = \{A, E\}$, where an *agent* (A) in-
 10 teracts with an *environment* (E). The agent
 11 $A = \{B, M\}$ is formed by two subsystems:
 12 *body* (B) and *mind* (M). The body is the in-
 13 termediary between mind and environment.
 14 Mind, body and environment can be each
 15 described by an abstract state space and an
 16 evolution function: $E = \{X_E, f_E\}$, $B = \{X_B, f_B\}$,
 17 $M = \{X_M, f_M\}$.

18 « 15 » These entities are interrelated
 19 dynamical systems. The environment con-
 20 tinually imposes a *situation* (s) on the agent,
 21 which responds through an *actuation* (a).
 22 The situation is given in function of the state
 23 of the environment, $s = f_s(x_E)$, and the actua-
 24 tion is defined according to the state of the
 25 body, $a = f_a(x_B)$. In the same way, the mind
 26 is continually receiving a perception signal
 27 coming from the body in function of its
 28 state, $p = f_p(x_B)$, and sending to the body a
 29 *control* signal (c), decided in function of the
 30 mind’s own internal state, $c = f_c(x_M)$. Part of
 31 the situation can be perceived by the mind
 32 through *external sensors* present in the body,
 33 while the mind can also control part of the
 34 actuation over the environment through *ex-
 35 ternal effectors* also present in the body. The
 36 interaction of the mind with the body takes
 37 place through *internal sensors* and *effectors*.
 38 The mind does not know *a priori* what sen-
 39 sors and effectors are internal or external.
 40 From the point of view of the mind, both
 41 body and environment are in some way ex-
 42 ternal, being part of an *exteriority* $W = \{B,$
 43 $E\}$, the world outside the mind. The com-
 44 plete CAES meta-architecture is presented
 45 in Figure 1.

46 « 16 » This configuration generates a
 47 kind of circularity, and defines each entity as
 48 a partially open dynamical system. The envi-
 49 ronment evolves in function of its own cur-
 50 rent state, but influenced also by the actua-
 51 tion coming from the agent, $x_E' = f_E(x_E, a)$.
 52 Similarly, the next body state is defined in
 53 function of the actual body state, but is influ-
 54 enced by both the situation coming from the
 55 environment and the control signal coming

column A

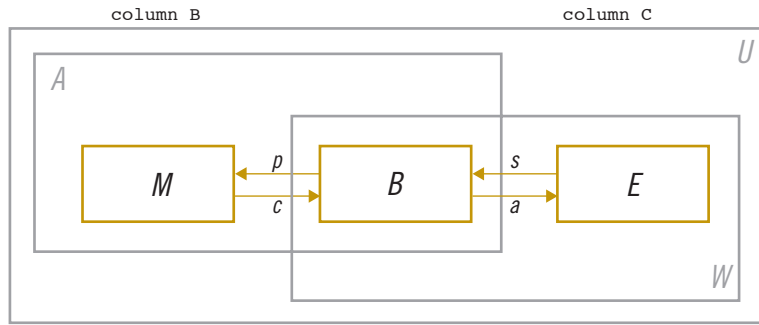


Figure 1: Body between mind and environment.

from the mind, $x_B' = f_B(x_B, s, c)$. It is the same
 for the mind, which continually changes its
 internal state (whatever that means) influ-
 15 enced by its perceptions, $x_M' = f_M(x_M, p)$.

16 « 17 » CAES is a meta-architecture
 because it does not define of what or how
 each system is made. Moreover, it does not
 17 constraint these systems as *stationary*.³ The
 environment as well as the body can change
 its respective set of rules and variables over
 time. The same applies for the mind, which
 needs to be non-stationary if we want to
 have some kind of learning or mental devel-
 opment. Such learning ability can be defined
 as a function $M' = f_\mu(M, x_M, p)$ that changes
 the mind’s own space of states (creating
 new concepts or representational signs) and
 rules (changing the policy of actions that is
 responsible for determining the control sig-
 18 nal) based on the experience (memories and
 immediate perceptions).

Representing ontological and experiential reality

19 « 18 » In our understanding, a con-
 20 structivist machine learning mechanism
 must be made using *model-based*⁴ methods.
 The agent constructs knowledge in order
 to understand its experience of interaction
 with the environment. Computationally,
 the learning problem can be divided into
 two parts: (a) the construction of the mod-

3| A dynamical system is stationary if the rules that define its evolution do not change over time.

4| In opposition to *model-free* methods, where an agent can dynamically optimize its behavior only based on the immediate experience.

el, and, based on it, (b) the definition of a
 20 policy of actions, which defines the agent’s
 21 subsequent behavior.

22 « 19 » When, for simplicity, we say that
 the agent constructs a *model of the world*, we
 need to specify that in fact it is the agent’s
 mind that constructs a model of an exteri-
 21 ority (the world outside the mind) to which
 the mind has access only through a limited
 22 sensorial interface. A model of the world is
 not a reproduction of the structure of an on-
 23 tological reality, but is a model of the agent’s
 24 *experiential history*.⁵ It is a model (and not a
 25 memory) because it does more than remem-
 26 ber the past interactions: the model aims to
 27 generalize a complete system to represent
 the whole external world based on the finite
 28 set of experiences.

29 « 20 » Frequently in the machine learn-
 30 ing literature, the relation between agent
 and environment is not clearly defined. Tra-
 31 ditionally, computer scientists do not make
 an explicit difference between the world as
 32 it is ontologically and the world represent-
 33 ed by the agent at the limits of its sensorial
 34 interface and history of interactions. They
 35 conceive the agent as acting and perceiving
 directly on the “real world,” and this can give
 36 rise to confusing architectures, where situ-
 37 ativity problems disappear by omission.

38 « 21 » We define the learning problem in
 the following terms: we cannot know what
 any “external reality” (the world outside the
 39 mind) consists of, but we suppose that it can
 40 be represented (for analytical purposes) as a
 41 *factored and partially observable Markovian
 42 decision process* (FPOMDP), where actions

5| The experiential history is the sequence of interactions (perceptions and actions) realized by the mind.

column B

column C

column A

1 and observations correspond to the control
2 and perception signals in the CAES archi-
3 tecture. The world model constructed by
4 the mind can be represented by a *factored*
5 *Markovian decision process* (FMDP) that
6 constitutes a kind of morphism of the first
7 one, constrained by the sensorial interface
8 limitations, as well as by the incompleteness
9 of the experience, but possibly enriched
10 with abstract variables created in order to
11 make the system more structured and intel-
12 ligible.

13 « 22 » In simulated systems, where both
14 agent and environment are programs run-
15 ning in a computer, an observer can have
16 access to the whole structure (mind, body,
17 environment and their interfaces of interac-
18 tion). In this particular case, it is possible
19 to analyze the factors that characterize the
20 experiential relation with that given reality.
21 The specificities of that relation, combined
22 with the intellectual and cognitive capaci-
23 ties of the agent, will determine the difficul-
24 ty of learning a *successful model*,⁶ and con-
25 sequently the agent's possibility to become
26 adapted to the environment.

MDP framework

27
28
29
30
31 « 23 » *Markovian decision processes*
32 (MDPs) and their extensions constitute
33 widely-used representations for modeling
34 *decision-making* and *planning* problems
35 (Feinberg & Shwartz 2002). An MDP is
36 typically represented as a discrete stochastic
37 *finite state machine* (Puterman 1994; Rivest
38 & Schapire 1994): at each time step the ma-
39 chine is in some state s ; the agent interacts
40 with the process by choosing some action a
41 to carry out; then the machine changes into
42 a new state s' and gives the agent a corre-
43 sponding reward r ; a given transition func-
44 tion δ defines the probabilities of the state
45 change according to s and a . The flow of an
46 MDP (the transition between states) de-
47 pends only on the system's current state and
48 on the action taken by the agent at the time.
49 After acting, the agent receives an evalua-

50
51 6| A model can be considered successful if
52 it allows the agent to make correct anticipations
53 for future interactions; it is not an evaluation of
54 the correspondence with ontological structures,
55 which remain inaccessible.

column A

column B

tive reward signal (positive or negative), ac-
cording to the chosen actions or the realized
state transition.

« 24 » Solving an MDP means find-
ing the optimal (or near-optimal) policy of
actions in order to maximize the rewards
received by the agent over time. When the
MDP parameters are completely known, in-
cluding the reward and the transition func-
tions, it can be mathematically solved by *dy-*
namic programming methods. When these
functions are unknown, the MDP can be
solved by *reinforcement learning* methods,
designed to learn a policy of actions on-line,
i.e., at the same time that the agent interacts
with the system, by incrementally estimat-
ing the utility of state-actions pairs and then
mapping situations to actions (Sutton &
Barto 1998).

« 25 » However, the MDP supposes that
the agent has complete information about
the state of the environment. A partially ob-
servable MDP (POMDP) (Singh et al. 2003;
Cassandra, Kaelbling & Littman 1998) is an
extension of the model that includes a set of
observations that is different from the set of
states. The underlying system state s cannot
be directly perceived by the agent, which has
access only to an observation o given by an
observation function γ . We can represent a
larger set of problems using POMDPs rather
than MDPs, but the methods for solving
them are computationally even more expen-
sive (Hauskrecht 2000).

« 26 » For a situated agent, this kind of
representation becomes inadequate because
it requires the complete enumeration of the
states, and the number of states increases
exponentially according to the number of
agent sensors (Bellman 1957). This is the
main bottleneck in the use of MDPs or
POMDPs: representing complex universes
entails an exponential increase in the state
space, and the problem quickly becomes in-
tractable.

Factoring the MDP states

« 27 » When a large MDP has a signifi-
cant internal structure, it can be modeled
compactly; the factorization of states is an
approach to exploit this characteristic (Bou-
tilier, Dearden & Goldszmidt 2000; Jonsson
& Barto 2005; Degris, Sigaud & Wuillemin
2006; Shani et al. 2008). In the factored rep-
resentation, a state is implicitly described

column B

column C

by an assignment to some set of state vari- 1
ables. Thus, a complete explicit state space 2
enumeration is avoided, and the system can 3
be described referring directly to its vari- 4
ables. The factorization of states enables the 5
system to be represented in a generalized 6
and compact way, even if the correspond- 7
ing MDP is exponentially large (Guestrin et 8
al. 2003). When the structure of the FMDP 9
is completely known, it is possible to find 10
good policies in an efficient way (Guestrin 11
et al. 2003). However, the research con- 12
cerning the discovery of the structure of 13
an underlying system from incomplete ob- 14
servation is still incipient (Degris & Sigaud 15
2010).

« 28 » An FPOMDP is an FMDP that 17
can represent partial observation (Guestrin, 18
Koller & Parr 2001; Hansen & Feng 2000; 19
Poupart & Boutilier 2004; Shani, Brafman & 20
Shimony 2005; Sim et al. 2008). An FPOM- 21
DP can be formally defined as a 4-tuple $\{X,$ 22
 $C, R, T\}$. The state space is factored and rep- 23
resented by a finite non-empty set of system 24
properties or variables $X = \{X_1, X_2, \dots, X_n\}$, 25
which is divided into two subsets, $X = P \cup H$, 26
where the subset P contains the observ- 27
able properties (those that can be accessed 28
through the agent's sensory perception), 29
and the subset H contains the hidden or 30
non-observable properties. Each property 31
 X_i is associated to a specified domain, which 32
defines the values the property can assume; 33
 $C = \{C_1, C_2, \dots, C_m\}$ represents the controlla- 34
ble variables, composing the agent actions; 35
 $R = \{R_1, R_2, \dots, R_k\}$ is a set of (factored) re- 36
ward functions, in the form $R_i: P_i \rightarrow \mathbb{R}$; and 37
 $T = \{T_1, T_2, \dots, T_n\}$ is a set of transformation 38
functions, such as $T_i: X \times C \rightarrow X_i$, defining 39
the system dynamics. Each transformation 40
function can be represented by a *dynamic* 41
Bayesian network, which is an acyclic, ori- 42
ented, two-layer graph. The first layer nodes 43
represent the environment state at time t , 44
and the second layer nodes represent the 45
next state, at $t+1$ (Boutilier, Dearden & 46
Goldszmidt 2000). A policy π is a mapping 47
 $X \rightarrow C$ where $\pi(x)$ defines the action to be 48
taken in x . The agent must learn a policy that 49
optimizes the average rewards received over 50
time, but it never sees the ontological state x , 51
only a perceptive situation p . 52

« 29 » When the agent is immersed in 53
a system represented as an FPOMDP, the 54
complete task for its anticipatory learning 55

column C



column B

column C

FILIPPO STUDZINSKI PEROTTO

received his doctorate degree from the University of Porto Alegre (UFRGS, Brazil) in cooperation with the University of Toulouse (INPT, France) in 2010, treating the subject of constructivist artificial intelligence models. He works with applied AI for industry, and participates in the constructivist AI research group of Toulouse.

Relation between ontological and experiential reality

« 30 » We distinguish four main factors that shape the relation between the agent's mind and the external world: *observability*, *complexity*, *determinism* and *controllability*.

« 31 » The *observability* factor (ω) indicates the degree of access that the agent has to the environment state through its sensorial perception. We can imagine this measure as being equivalent to the proportion of observable variables in the whole system in relation to the total number of variables. If the state of the environment can be represented by n bits of information and the state of the sensors affected by that world state can be represented by m bits, the observability factor ω is the proportion of m over n , where $0 \leq \omega \leq 1$. Considering an FPOMDP composed of binary variables, $m = |P|$ and $n = |X|$.

« 32 » If $\omega = 1$, the environment is said to be completely observable, which means that the agent has sensors to observe directly all the properties of the environment. In this case there is no perceptual confu-

sion, and the agent always knows the current state. When $\omega < 1$, the environment is said partially observable. The lower ω is, the higher the proportion of hidden dimensions of the environment is in relation to the agent's perception. When ω is close to 0, the agent is no longer able to identify the current situation only in terms of its perception.

« 33 » The *complexity* factor (φ) is related to the rules that define the world dynamics, indicating how intelligible the environment transformations can be for the agent. The complexity can be measured as the average amount of information needed to define the evolution of one bit in the world state. In a highly structured world, it is possible to model precise causes for each transformation; in other words, the evolution of one variable of the system depends on only a few other relevant variables. In contrast, in an unstructured world there is too much interdependence between the variables to determine the evolution of the system.

« 34 » The difficulty for the agent in constructing a model is related to the complexity of the world dynamics. A less complex world can be more easily structured by intelligence. A low level of complexity means that the information about the dynamics of the environment is concentrated in the variables. It indicates the average amount of relevant variables necessary to describe each transformation. When φ is small, the rules that govern the dynamics of the whole system have few parameters. It is a kind of thermometer indicating how easy is to model causality between events. In contrast, a higher level of complexity (rising to n) indicates that the information about the dynamics is sparsely distributed over

all the set of variables, and in this case the agent needs to describe the transformations in function of almost all the variables.

« 35 » The *determinism* factor (∂) is equivalent to the proportion of deterministic transformations in relation to the total number of transformations. In the completely non-deterministic case ($\partial = 0$), all transformation functions (of every probability) need to be represented in terms of probabilities. On the other hand, in the completely deterministic case ($\partial = 1$), every transformation is deterministic. An environment is said partially deterministic if it is situated between these two extremities ($0 < \partial < 1$) presenting both deterministic and stochastic transformations.

« 36 » Observability and determinism are dependent factors. Partially observable environments can present some determinant variables to a good world model that cannot be directly perceived by the agent sensors. Such environments can appear arbitrarily complex and non-deterministic on the surface, but they can actually be deterministic and predictable with respect to unobservable underlying elements (Holmes & Isbell 2006). In other words, an ontologically deterministic world can be experienced as non-deterministic. The more an agent has sensors to perceive complex elements and phenomena, the more the environment will appear deterministic to it.

« 37 » Finally, the *controllability* factor (κ) represents the proportion of variables whose dynamics are influenced by the agent's actions, within the total number of variables in the system. The controllability factor affects the difficulty of learning because it determines the capacity of the agent to experiment actively.

column A

column B

column C

CALM

column A

column B

column C

1 CALM
2
3 «38» The *constructivist anticipatory*
4 *learning mechanism* (CALM), detailed in
5 (Perotto 2010), is a mechanism that enables
6 an agent to learn the structure of an un-
7 known environment in which it is situated
8 through observation and experimentation,
9 creating an anticipatory model of the world.
10 CALM operates the learning process in an
11 *active* and *incremental* way. There is no sepa-
12 rated previous training time: the agent has a
13 single uninterrupted interactive experience
14 within the system; it needs to perform and
15 learn at the same time.

16 «39» The task becomes harder because
17 the environment is only partially observable
18 and partially deterministic, from the point of
19 view of the agent, constituting an FPOMDP.
20 In this case, the agent has perceptive infor-
21 mation from a subset of sensory variables,
22 but the system dynamics also depends on an-
23 other subset of hidden variables. To be able
24 to create a consistent world model, the agent
25 needs, beyond discovering the regularities
26 of the phenomena, also to create abstract
27 variables in order to take into account non-
28 observable conditions that are necessary to
29 understand the system's evolution. In other
30 words, learning a model of the world is more
31 than describing the environment dynam-
32 ics (the rules that can explain and anticipate
33 the observed transformations), it is also dis-
34 covering the existence of hidden properties
35 (once they influence the evolution of the ob-
36 servable ones) and, finally, finding a way to
37 deduce the values of these hidden properties.

38 «40» The system as a whole is in fact an
39 FPOMDP, but CALM is designed to discover
40 the existence of non-observable properties,
41 integrating them in its anticipatory model.
42 In this way CALM can infer a structure to
43 represent the dynamics of the system in
44 the form of an FMDP (if the agent can suc-
45 cessfully discover and describe the hidden
46 properties of the FPOMDP that it is dealing
47 with, then the world becomes treatable as an
48 FMDP because the hidden variables become
49 known). There are some algorithms able to
50 calculate efficiently the optimal (or near-
51 optimal) policy, when the FMDP is given
52 (Guestrin et al. 2003). The algorithm to cal-
53 culate the policy of actions used by CALM is
54 similar to that presented by Degris, Sigaud &
55 Wuillemin (2006). However, the main chal-

column A

lenge is to discover the structure of the prob-
lem based on the on-line observation. CALM
does it using representations and strategies
inspired by Drescher (1991).

Representing predictive knowledge by schemas

«41» CALM tries to reconstruct, by ex-
perience, each system transformation func-
tion T_i , representing it by an *anticipatory tree*.
Each anticipatory tree is composed of pieces
of predictive knowledge called *schemas*; each
schema represents some perceived regularity
occurring in the environment by associating
context (sensory and abstract), actions and
expected results (anticipations).

«42» One important strategy for deal-
ing with complexity is finding what is impor-
tant to anticipate. At the beginning, the only
interesting variables are those associated to
positive or negative affective values. Staying
focused on these variables avoids wasting
energy by creating models that anticipate
other non-important variables. Gradually,
the variables needed to anticipate the evo-
lution of some important variable (relation
of causality) are also considered important,
and the mechanism will seek to model their
transformation function too.

«43» A schema is composed of three
vectors, in the form

$$\Xi = \{context \wedge action \rightarrow result\}$$

denoting a kind of predictive rule. The *con-*
text vector has their elements linked both
with the agent sensors and with the abstract
variables. These abstract variables are rep-
resented by (mentally created) “synthetic
elements” not linked to any sensor but refer-
ring to non-sensory properties of the uni-
verse, the existence of which is inferred by
the mechanism. The action vector is linked
with the agent effectors. Context and ac-
tion vectors can represent sets of equivalent
situations or actions, by generalization. The
result vector represents the value expected
for some variable in the next time, after ex-
ecuting the given action in the given context.
Each element vector can assume any value in
a discrete interval defined by the respective
variable domain.

«44» Some elements in these vectors
can take an “undefined value.” For example,
an element linked with a binary sensor must
have one of three values: true, false or unde-

column B

defined (represented, respectively, by “1”, “0”
and “#”). The undefined value generalizes the
schema because it allows some properties to
be ignored in order to represent a set of situa-
tions. The learning process happens through
the refinement of the set of schemas. After
each experienced situation, CALM updates
a generalized episodic memory, then checks
whether the result (context perceived at the
instant following the action) conforms to the
expected result of the activated schema. If
the anticipation fails, the error between the
result and the expectation serves as param-
eter to correct the model. The context and
action vectors are gradually specialized by
differentiation, adding each time a new rel-
evant feature to identify the situation class
more precisely.

«45» The use of undefined values
makes it possible to construct an *anticipatory*
tree. Each node in that tree is a schema, and
relations of generalization and specialization
guide its topology (quite similar to decision
trees or discrimination trees). The root node
represents the most generalized situation,
in which the context and action vectors are
completely undefined. Each level added to
the tree represents the specialization of one
element, where each branch replaces the
undefined (generalized) value with one dif-
ferent possible defined value. This specializa-
tion occurs either in the context vector or in
the action vector. In this way, CALM divides
the state space according to the different
expected results, grouping contexts and ac-
tions with their respective transformations.
The tree evolves during the agent's life, and
is used by the agent, even if the tree is still
under construction, to take its decisions, and
in consequence, to define its behavior. The
structure of a schema (the elementary piece
of knowledge of an anticipatory tree) is pre-
sented in Figure 2.

«46» The context in which the agent
is at a given moment (perceived through its
sensors) is applied in the tree, exciting all
the schemas that have a compatible context
vector. This process defines a set of excited
schemas, each one suggesting a different ac-
tion to take in the given situation. CALM
will choose one action to activate and will
perform it through the agent's effectors. The
algorithm always chooses the compatible
schema that has the most specific context,
called decider schema, which is the leaf of a

column C

column A

column B

column C

1 differentiated branch. This decision is taken
 2 based on the calculated utility of each possi-
 3 ble choice. There are two kinds of utility: the
 4 first estimates the discounted sum of rewards
 5 in the future following the policy, the second
 6 measures the exploration benefits. The util-
 7 ity value used to take the decision depends
 8 on the circumstantial agent strategy (exploit-
 9 ing or exploring). The mechanism also has a
 10 kind of generalized episodic memory, which
 11 represents (in a compact form) the specific
 12 and real situations experienced in the past,
 13 preserving the information necessary to cor-
 14 rectly construct the tree.

15
 16 **Anticipatory tree construction**

17 « 47 » The learning process happens
 18 through the refinement of the set of sche-
 19 mas. At each given moment in the time,
 20 the set of schemas of our agent, gradually
 21 constructed by the mechanism, is assumed
 22 to be coherent with all the past experience,
 23 describing in an organized way the regular
 24 phenomena observed during the interaction
 25 with the universe. To do so, the mechanism
 26 must have a memory of the past situations,
 27 but this memory can be neither too precise
 28 (because remembering all the experienced
 29 episodes would require a nonviable amount
 30 of space) nor too simple (because the lack
 31 of information would make it impossible to
 32 revise the model if there was contradiction
 33 with new disequilibrating observations).
 34 The implementation of a feasible episodic
 35 memory is not evident; it can be very expen-
 36 sive if we try to stock too much information
 37 coming from the sensory flow. However, us-
 38 ing some strong but well-chosen restrictions
 39 (such as limiting the dependency analysis
 40 between variables), and using a generalized
 41 and structured representation of the past ex-
 42 perience, it becomes computationally viable.

43 « 48 » After each experienced situation,
 44 CALM actualizes the generalized episodic
 45 memory and checks whether the result (con-
 46 text perceived at the instant following the
 47 action) is in conformity to the expectation
 48 of the activated schema in the anticipatory
 49 tree. If the anticipation fails, the error be-
 50 tween the result and the expectation serves
 51 as a parameter for correcting the model. In
 52 the anticipatory tree topology, the context
 53 and action vectors are taken together. This
 54 concatenated vector identifies the node
 55 in the tree. It can be expanded following a

column A

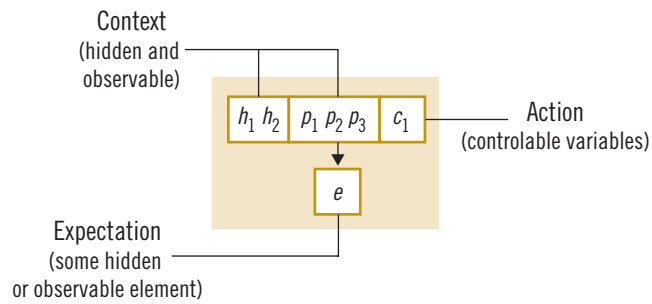


Figure 2: The anticipatory tree. Each node is a schema composed of three vectors: context, action and expected result; the leaf nodes are decider schemas.

top-down strategy: the initial tree contains a unique schema, with completely generalized context and action, and it is gradually specialized by differentiation, adding new relevant features to identify more precisely the category of equivalent situations, which entails the creation of new branches in the tree where the context and action vectors are each time more defined. In well-structured universes, the shorter way is starting with an empty vector and searching for the probably small set of features relevant to distinguish the important situations, rather than starting with a full vector and having to waste energy eliminating a lot of useless elements. Selecting the right set of relevant features to represent some given concept is a well-known problem in AI, and the solution is not easy, even using approximated approaches. To do this, CALM adopts a forward greedy selection (Blum & Langley 1997), using the data registered in the generalized episodic memory.

« 49 » The *expected result* vector can be seen as a label in each decider schema, anticipating how the world changes when the schema is activated. Initially, all different expectations are considered as different classes, and they are gradually generalized and integrated with others. The agent has two alternatives when the expectation fails. In a way that makes the knowledge compatible with the experience, the first alternative is to try to divide the scope of the schema, creating new schemas with more specialized contexts. Sometimes this is not possible and then the schema's expectation is reduced. In the expected result vector, “#” means that

column B

the element is not deterministically predict-
 able. Another symbol can be used to rep-
 resent some special situations, in order to
 reduce the number of schemas; this is the
 symbol “=”, used to indicate that the value
 of the expected element will not be changed.

« 50 » Three basic methods compose
 the CALM learning function, namely: *dif-*
ferentiation, *adjustment*, and *integration*.
 Differentiation is a necessary mechanism
 because a schema responsible for a too
 general context cannot often make precise
 anticipations. If a general schema does not
 work well, the mechanism divides it into
 new schemas, differentiating them by some
 element of the context or action vector. In
 fact, the differentiation method takes an un-
 stable decider schema and changes it into
 a two level sub-tree. The parent schema in
 this sub-tree preserves the context of the
 original schema. The children, which are the
 new decider schemas, have context vectors
 that are a little more specialized than those
 of their parent. They attribute a value to
 some undefined element, dividing the scope
 of the original schema. Each one of these
 new deciders engages itself in a part of the
 domain. In this way, the previous correct
 knowledge remains preserved, distributed
 in the new schemas, and the discordant situ-
 ation is isolated and treated only in its spe-
 cific context. Differentiation is the method
 responsible for making the anticipatory tree
 expand. Each level of the tree represents the
 introduction of some new constraint. The
 algorithm needs to choose what will be the
 differentiator element, which could be from
 either the context vector or the action vec-

column C




Figure 3: Differentiation method example: (a) a real experimented situation (with five variables) and executed action (one variable); (b) activated schema (with compatible context, action, and expectation); (c) associated episodic memory (representation of real situations where the scheme has been activated, in this case representing no interdependencies between variables); (d) real observed result, after the execution of the action; (e) sub-tree generated by differentiation in order to compensate the divergence observed between expectation and result.




Figure 4: Adjustment method example: (a) a real experimented situation and action; (b) activated schema; (c) associated episodic memory; (d) real observed result; (e) schema expectation reduction after adjustment.




Figure 5: Integration method: (a) sub-tree after an adjustment; (b) an integrated schema substitutes the sub-tree.




Figure 6: Synthetic element creation method: (e) incremented context and expectation vectors, and differentiation using a synthetic element.

column C
tor. This differentiator needs to separate the situation responsible for the disequilibrium from the others, and the algorithm chooses it by calculating the information gain, and considering a limited (parametrized) range of interdependencies between variables. Figure 3 illustrates the differentiation process.

« 51 » When some schema fails and it is not possible to differentiate it in any way, then CALM executes the adjustment method. This method reduces the expectations of an unstable decider schema in order to make it reliable again. The algorithm simply compares the activated schema's expectation and the real result perceived by the agent after the application of the schema, setting the incompatible expectation elements to the undefined value (“#”). The adjustment method changes the schema's expectation (and consequently the anticipation predicted by the schema). Figure 4 illustrates this.

« 52 » In this way, the schema expectation can change (and consequently the class of the situation represented by the schema), and the tree maintenance mechanism needs to be able to reorganize the tree when this change occurs. Therefore, successive adjustments in the expectations of various schemas can reveal unnecessary differentiations. When CALM finds a group of schemas with similar expectations for approaching different contexts, the integration method comes into action, trying to join these schemas by searching for some unnecessary common differentiator element and eliminating it. The method operates as shown in Figure 5.

Dealing with the unobservable

« 53 » When CALM reduces the expectation of a given schema by adjustment, it assumes that there is no deterministic regularity following the represented situation in relation to these incoherent elements, and that the related transformation is unpredictable. However, sometimes a prediction error can be explained by considering the existence of some abstract or hidden property in the environment, which could be useful to differentiate an ambiguous situation but which is not directly perceived by the agent sensors. So, before adjusting, CALM assumes the existence of a non-sensory property in the environment, which will be represented as a synthetic element. When a new synthetic element is created, it is included as

column A

1 a new term in the context and expectation
2 vectors of the schemas. The use of synthetic
3 elements assumes the existence of some-
4 thing beyond the sensory perception, which
5 can be useful to explain non-equilibrated
6 situations. They have the function of ampli-
7 fying the differentiation possibilities.

8 « 54 » In this way, when dealing with
9 partially observable environments, CALM
10 has two additional challenges: (a) inferring
11 the existence of unobservable properties,
12 which it will represent by synthetic ele-
13 ments, and (b) including these new elements
14 into its predictive model. A good strategy
15 for doing this is to look at the historical in-
16 formation.

17 « 55 » CALM introduces a method
18 called *abstract differentiation*. When a
19 schema fails in its prediction, and when it
20 is not possible to differentiate it by the cur-
21 rent set of considered properties, then a
22 new Boolean synthetic element is created,
23 enlarging the context and expectation vec-
24 tors. Immediately, this element is used to
25 differentiate the incoherent situation from
26 the others. The method attributes arbitrary
27 values to this element in each differentiated
28 schema. These values represent the presence
29 or absence of some non-observable condi-
30 tion, necessary to determine the correct pre-
31 diction in the given situation. The method
32 is illustrated in Figure 6, where the new ele-
33 ments are represented by card suits.

34 « 56 » Once a synthetic element is cre-
35 ated, it can be used in subsequent differen-
36 tiations. A new synthetic element will be
37 created only if the existing ones are already
38 saturated. To avoid the problem of creat-
39 ing infinite new synthetic elements, CALM
40 can do this only up to a determined limit,
41 after which it considers that the problematic
42 anticipation is not deterministically pre-
43 dictable, undefining the expectation in the
44 related schemas by adjustment. Figure 7 il-
45 lustrates the idea behind synthetic element
46 creation.

47 « 57 » The synthetic element is not as-
48 sociated to any sensory perception. Con-
49 sequently, its value cannot be observed.
50 This fact can place the agent in ambiguous
51 situations, where it does not know whether
52 some relevant but non-observable condition
53 (represented by this element) is present or
54 absent. Initially, the value of a synthetic ele-
55 ment is verified *a posteriori* (i.e., after the ex-

column A

column B

column C

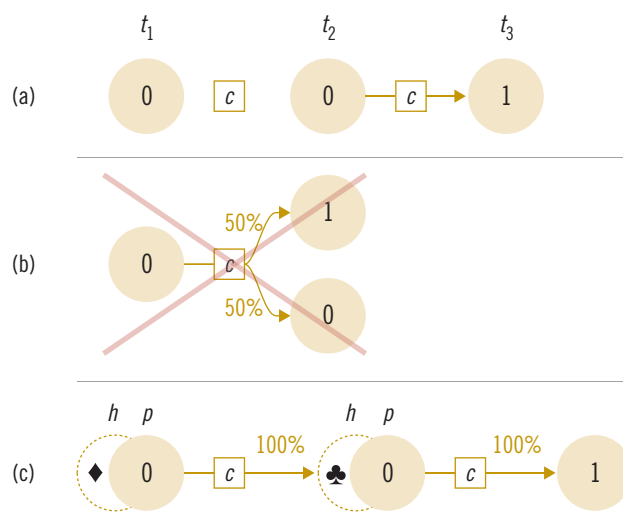


Figure 7: Discovering the existence of non-observable properties in: (a) a real experienced sequence; (b) what CALM does not do (the attribution of a probability); (c) the creation of a synthetic element in order to explain the observed difference




Figure 8: Predicting the dynamics of a non-observable property in: (a) a real experienced sequence; (b) the use of a synthetic element to explain the logic behind the observed transformations.

ecution of the action in an ambiguous situ-
ation). Once the action is executed and the
following result is verified, then the agent
can rewind and deduce the situation really
faced in the past instant (disambiguated).
Discovering the value of a synthetic element
after the circumstance where this informa-
tion was needed can seem useless, but in fact
this delayed deduction gives information to
another method called *abstract anticipation*.
If the non-observable property represented
by this synthetic element has a regular dy-
namics, then the mechanism can propagate
the deduced value back to the schema acti-
vated in the immediately previous instant.
The deduced synthetic element value will be
included as a new anticipation in the previ-

ous activated schema. Figure 8 shows how
this new element can be included in the pre-
dictive model.

« 58 » For example (complementing
Figure 8), in time t_1 CALM activates a sche-
ma $\Xi_1 = \{\#1 \wedge c \rightarrow \#0\}$, where the context and
expectation are composed of two elements
(the first one synthetic and the second one
perceptive) and one action. Suppose that the
schema succeeds and, as predicted, the next
observation is “0”. The problem is that the
next situation “#0” is ambiguous because it
excites both the schemas, $\Xi_2 = \{\#0 \wedge c \rightarrow \#0\}$
and $\Xi_3 = \{\#0 \wedge c \rightarrow \#1\}$. At this time, the
mechanism cannot know the value of the
synthetic element, crucial to determining
the real situation. Suppose that, anyway, the

column B

column C

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55




Figure 9: The hyper-flip problem.

mechanism decides to execute the action “c” in time t_2 , and this is followed by the sensory perception “0” in t_3 . Only now, in t_3 , the agent can deduce that the situation really dealt with in t_2 was “ $\spadesuit 0$ ”, and it can include this information in the schema activated in t_0 , in the form $\Xi_1 = \{\#1 \wedge c \rightarrow \spadesuit 0\}$.

Experimental results

« 59 » To exemplify the functioning of the proposed method, we will use the hyper-flip problem, and extension of the problem used by Satinder Singh et al. (2003) and Michael Holmes & Charles Isbell (2006). It consists of an agent who lives in a two-state universe. It has 3 actions (l , r , u) and 2 perceptions (0, 1). The agent does not have any direct perception of the underlying current state. It sees “1” when the state changes horizontally, and “0” otherwise. Action “u” changes the state vertically, action “l” causes the deterministic transition to the left state,

and action “r” causes the deterministic transition to the right state. The flip problem is showed as a state machine in Figure 9.

« 60 » CALM is able to solve this problem. First, the mechanism tries to predict the next observation in function of its action and current observation. However, it quickly discovers that the perceptive observation is not useful to the model, and that there is insufficient information to make correct anticipations. So, it creates a new synthetic element that will be able to represent the underlying left (\clubsuit) and right (\spadesuit) states. Figure 10 shows the final solution. It is interesting to note that the constructed world model (with its 3 variables) is not a copy of the ontological structure of the problem (a machine with 4 states).

« 61 » In order to test the robustness of the mechanism, a hundred new observable variables have been inserted in the hyper-flip problem for a second scenario. These new variables present random transformation functions and do not influence the evolution of the original observation. The result is that the mechanism is not affected in its capacity to solve the problem (it finds the same solution as that previously indicated). The time of learning increases in a linear order with the addition of irrelevant variables.

Related work

« 62 » CALM is an original mechanism that enables an agent to create incrementally a model of an experience during the course of its interaction with the universe. The pioneer work on constructivist AI was pre-

sented by Drescher (1991). He proposed the first constructivist agent architecture, which learns a world model by an exhaustive statistical analysis of the correlation between all the context elements observed before each action, combined with all resulting transformations. Drescher has also suggested the need to discover hidden properties by creating “synthetic items.”

« 63 » The schema mechanism represents a strongly coherent model. However, there are no theoretical guarantees of convergence. Another restriction is the computational cost of the kind of operations used in the algorithm. The need for space and time resources increases exponentially with the problem size. Nevertheless, some other researchers have presented alternative models inspired by Drescher, such as Yavuz & Davenport (1997), Morrison, Oates & King (2001), Chaput (2004), and Holmes & Isbell (2005), always based on the search for statistically observed regularities.

« 64 » CALM differs from these previous works because we limit the problem to the discovery of deterministic regularities (even in partially deterministic environments). In this way, we can implement direct induction methods in the agent learning mechanism. This approach presents a low computational cost, and it allows the agent to learn incrementally and find high-level regularities. For that, we have been inspired by Holmes & Isbell (2006), who used the notion of the state signature as a historical identifier of the states to develop the idea of learning anticipations through the analysis of relevant pieces of history.

« 65 » With the emergence of the factored MDP model, some important works have been realized to create algorithms designed to discover the structure of the system (Degris, Sigaud & Wuillemin 2006; Degris & Sigaud 2010; Strehl, Diuk & Littman 2007; Jonsson & Barto 2005). However CALM, as far as we know, is the only one to merge the induction of synthetic elements to represent the non-observable variables in an FPOMDP.

« 66 » Another originality of CALM is the use, in such learning problems, of a generalized episodic memory associated to the search for important variables (related to affective values or relevant to anticipate the evolution of other important variables).

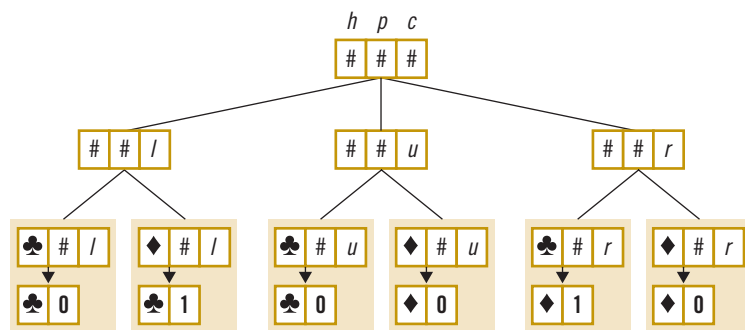


Figure 10: Final schematic tree for solving the flip problem. The vector represents synthetic elements (h), perceptible elements (p) and actions (c). The decider schemas show the expectations.

column A

Conclusion

« 67 » The CALM mechanism can provide autonomous adaptive capabilities to an agent because it is able to construct knowledge incrementally to represent the deterministic regularities observed during its interaction with the environment, even in partially deterministic universes.

« 68 » CALM is able to deal with partially observable environments, detecting high-level regularities. The strategy is the induction and prediction of unobservable properties, represented by synthetic elements.

« 69 » Synthetic elements enable the agent to step beyond the limit of instantaneous and sensorimotor regularities. In the agent's mind, synthetic elements can represent

column B

sent three kinds of "unobservable things": (a) hidden properties in partially observed worlds, or sub-environment identifiers in discrete non-stationary worlds; (b) markers to necessary steps in a sequence of actions, or to different possible agent points of view; and (c), abstract properties, which do not exist properly, but which are powerful and useful tools for the agent, enabling it to organize the universe into higher levels.

« 70 » With these capabilities, CALM is able to step beyond sensorial perception, constructing more abstract terms to represent the universe and to "understand" its own reality in more complex levels. CALM can be very effective for constructing models in partially but highly deterministic ($1 > \partial \gg 0$) and partially but highly observable ($1 > \omega \gg 0$) environments, and when

column C

the transformation functions have well-structured causal dependencies ($0 < \varphi \ll n$).

« 71 » Currently, we are improving CALM to enable it to form action sequences by chaining schemas. It will allow the creation of composed actions and plans. The next research steps include: formally demonstrating the mechanism's robustness and correctness; making comparisons between CALM and related solutions proposed by other researchers; and analyzing the mechanism's performance when facing more complex problems. Future works could include the extension of CALM to deal with non-deterministic regularities, noisy environments and continuous domains.

RECEIVED: 12 JULY 2013

ACCEPTED: 19 SEPTEMBER 2013

Open Peer Commentaries on Filippo Studzinski Perotto's "Computational Constructivist Model"



To Bridge the Gap between Sensorimotor and Higher Levels, AI Will Need Help from Psychology

Frank Guerin

University of Aberdeen, UK

f.guerin/at/abdn.ac.uk

> Upshot • Constructivist theory gives a nice high-level account of how knowledge can be autonomously developed by an agent interacting with an environment, but it fails to detail the mechanisms needed to bridge the gap between low levels of sensorimotor data and higher levels of cognition. AI workers are trying to bridge this gap, using task-

column A

specific engineering approaches, without any principled theory to guide them; they could use help from psychologists.

« 1 » The formulation of the problem as it appears in the abstract of Filippo Perotto's article packs in a lot of information that merits discussion:

“The constructivist conception of intelligence is very powerful for explaining how cognitive development takes place. However, until now, no computational model has successfully demonstrated the underlying mechanisms necessary to realize it. In other words, the artificial intelligence (AI) community has not been able to give rise to a system that convincingly implements the principles of intelligence as postulated by constructivism, and that is also capable of dealing with complex environments.”

column B

« 2 » This suggests that the psychologists have succeeded in explaining how cognitive development takes place and that the AI community has failed in its job to implement these "principles of intelligence." However, I would throw the problem back at the psychologists. I think that significant work is still needed at the level of theoretical psychology before we have something close to a proper explanation of how cognitive development takes place. Psychological explanations are for the most part vague and woolly; they do not elucidate the mechanisms underlying development (Jean Piaget's theory being a good example). Furthermore, Piaget's theory is at times even at odds with experimental psychology. It may be many, many years before we have a suitably detailed theory from the psychologists that is consistent with the evidence from experi-

column C

column A

1 ments. Until that time, one could argue that
2 the “principles of intelligence as postulated
3 by constructivism” are implemented very
4 well by existing AI systems. Gary Drescher,
5 for example, did implement the basic princi-
6 ples of constructivism, but “it could never be
7 used to solve significant applied problems,”
8 because the techniques do not scale up to
9 systems with large numbers of inputs and
10 degrees of freedom. However, Piaget did not
11 give us any idea of how to deal with these
12 issues, so one could lay the blame on him.

13 « 3 » To quote from the abstract again,
14 “there is a large distance between the de-
15 scriptions of the intelligence made by con-
16 structivist theories and the mechanisms
17 that currently exist.” If we consider Piaget’s
18 theory, and Drescher’s system or the CALM
19 system, I am not sure that there is such a
20 large distance. Piaget’s descriptions of as-
21 similation and accommodation are so all
22 encompassing and so lacking in detail that
23 it seems to me that Drescher’s system or
24 the CALM system constitute perfectly good
25 implementations. Psychology tends to leave
26 mechanisms very underspecified.

27 « 4 » To quote again from the article’s
28 abstract: “...and that is also capable of deal-
29 ing with complex environments.” Here is
30 perhaps the essence of the problem. When
31 you start building an actual AI system that
32 has to interact with the world, you face a
33 daunting task of dealing with a complex en-
34 vironment. It seems that AI is being saddled
35 with the burden of not only implementing
36 the high-level theory, but also making sure
37 it can deal with complex environments. The
38 “complex environments” problem needs to
39 be thrown back at the psychologists. The
40 history of AI has shown that a theory of
41 cognition that works at a high abstract level
42 but cannot account for the interface to the
43 sensorimotor level is not much of a theory
44 of cognition at all. The devil is in the detail.
45 There are many writers who convincingly
46 show how high-level cognition is very much
47 grounded in our sensorimotor intelligence
48 (e.g., Barsalou 2008; Byrne 2005; Bril, Roux
49 & Dietrich 2005). Psychological theories
50 tend to overlook the need for complex
51 mechanisms to bridge the gap between the
52 sensorimotor level and high-level cognition.
53 Psychologists may need to become compu-
54 ter scientists to some extent, so that they
55 have an appreciation of the computational

column A

column B

problems involved and the need for them
to describe mechanisms to account for how
humans successfully solve these.

« 5 » On the positive side, there are
some works in cognitive science that are
beginning to attempt to address the issue
of providing some theoretical framework
to account for how a sensorimotor level can
connect with higher levels of cognition: for
example, the multi-layered cognitive system
of Bipin Indurkha (1992, Chapter 5).

« 6 » For the CALM system itself, I feel
the article has all the correct ideas from a
philosophical and psychological point of
view, e.g., about the agent constructing its
own symbolic structures and not having
access to the “ontological reality.” However,
if we are to evaluate it as a candidate for a
“general artificial intelligence mechanism
that learns like humans do” (first sentence
of abstract), then it might suffer the same
shortcomings as Drescher’s work, i.e., “it
could never be used to solve significant ap-
plied problems.” For example, if the context
were to be the visual input from two stereo
cameras delivering a few million pixels in 24
bit colour at thirty frames per second and the
system is trying to predict the consequences
of actions, in the complexity of an everyday
setting, in this visual stream, it might not be
feasible to use each bit of input as a CALM
variable. One could, of course, propose
to hook the CALM system up to a higher-
level abstracted version of the visual input,
but then one runs into the issues of where
to make the cut-off between what the core
CALM system sees and what is the respon-
sibility of other abstraction mechanisms. If
the cut-off is at the wrong place, then one
runs into classical AI problems of (a) having
a core cognition that makes unreasonable as-
sumptions about how accurately it can inter-
face with the world or (b) having a prespec-
ified worldview imposed by the provided
abstractions (see Brooks 1991 or Stoytchev
2009 for problems with this). There does not
seem to be any clear theory from psychol-
ogy to guide us on how to connect the sen-
sorimotor level with some higher levels. AI
does have various different applied systems
that successfully make a connection from
high-level symbols to perception and action
in complex settings: for example, robots that
perform everyday tasks (Beetz et al. 2010).
However, each applied AI system tends to be

column B

column C

specialised and optimised for one particular
task. None could claim to be a reasonable
model of general human cognition, nor do
they attempt to be. This is really a job for the
psychologists.

Frank Guerin obtained his Ph.D. degree from Imperial
College, London, in 2002. Since August 2003, he has
been a Lecturer in Computing Science at the University
of Aberdeen. He is interested in understanding the core
of cognition in computational terms. He has focused
on understanding infant cognitive development,
as a first step to understanding later cognition.

RECEIVED: 11 OCTOBER 2013

ACCEPTED: 18 OCTOBER 2013

Environments Are Typically Continuous and Noisy

Martin V. Butz

University of Tübingen, Germany
martin.butz/at/uni-tuebingen.de

> **Upshot** • The schema system present-
ed in the target article suffers from prob-
lems that had been acknowledged more
than ten years ago. The main point is that
our world is neither deterministic nor
symbolic. Sensory as well as motor noise
is ubiquitous in our environment. Sym-
bols do not exist a priori but need to be
grounded within our continuous world.
In conclusion, I suggest that research on
schema-learning systems should tackle
small but real-world, continuous, and
noisy problem domains.

Heuristic learning principles are not enough

« 1 » About 15 years ago, I began work-
ing together with Wolfgang Stolzmann and
Joachim Hoffmann on the development of
anticipatory classifier systems (Stolzmann
2000). We attempted to tackle the funda-
mental problems of learning a cognitive
model in well-structured environments,
implementing contextual rule differentia-
tion, rule adjustment, and rule integration
mechanisms. With iterative improvements
and additions, the ACS2 system was devel-
oped. ACS2 combines a heuristic rule differ-

column C

column A

1 entiation and specialization mechanism that
2 is based on Hoffmann's cognitive learning
3 principle, termed "anticipatory behavioral
4 control" (Hoffmann 2003), with a gener-
5 alization mechanism that is implemented
6 by a steady-state evolutionary algorithm in
7 ACS2. In my book on *Anticipatory Learning*
8 *Classifier Systems* (Butz 2002), I summarized
9 the capabilities of the developed system as
10 well as the fundamental challenges.

11 « 2 » While the fundamental challenges
12 included the problem of *partially observ-*
13 *able Markov decision processes* (POMDPs),
14 I had also acknowledged that "essentially
15 any characteristic in an environment that
16 causes the deterministic perceptual causal-
17 ity to become probabilistic or noisy causes
18 difficulties" (Butz 2002: 127). I fear that the
19 algorithm presented in the target article suf-
20 fers similar difficulties. That is, while it may
21 be able to solve the tackled, small POMDP
22 problem, it is very doubtful that the heuris-
23 tic learning mechanism put forward is able
24 to produce similarly good solutions in noisy,
25 continuous environments.

26 « 3 » Is this a concern for the construc-
27 tivist community? In the following I will
28 argue that it is indeed a severe concern and
29 I propose that the community should focus
30 their research efforts on working with sys-
31 tems that experience noisy, continuous envi-
32 ronments rather than symbolic ones.

Noisy experiences

34 « 4 » Our Western world seems domi-
35 nated by symbolic knowledge, and so we
36 tend to forget our actual, natural environ-
37 ment. In this environment, our perceptions
38 are typically continuous and noisy, and
39 manipulations of and interactions with the
40 environment sometimes fail by nearly pure
41 chance. How can we live in this messy, non-
42 deterministic environment with all its cave-
43 ats? How can we learn a useful world model
44 with which we can manipulate and interact
45 with the environment purposefully?

46 « 5 » Various evidences suggest that our
47 mind constructs predictive models about
48 the consequences of body-environment in-
49 teractions (Butz 2008). Even in the simplest
50 cases, a certain form of causality is present
51 during such interactions. Thus, learning
52 about condition-action-effect contingen-
53 cies is possible and such knowledge is useful
54 when striving for a particular effect. How-

column A

column B

ever, a full-blown model of the environment
with all its inherent contingencies is too
large to grasp. Thus, as the author also sug-
gests in §42, most likely the models we learn
need to focus on those aspects that are rel-
evant for us, that is, those that are associated
with positive or negative affective values.

« 6 » However, these models need to
be functional in noisy, continuous environ-
ments. Thus, conditions, actions, and ef-
fects are initially not symbolic but consist of
contextual subspaces, motor primitives, and
local perceptual changes. According to Law-
rence Barsalou (1999) and others, we learn
our symbol processing capabilities during
our lives, grounding these capabilities in our
perceptual, noisy, and continuous experi-
ences.

« 7 » Most schema-oriented learning
systems, such as the one proposed in the
target article, have not managed to devel-
op symbol systems in a noisy, continuous
realm. Schema learning systems up until
now have stuck to symbol manipulation
problems, such as the admittedly tricky
hyper-flip problem. But are these problems
constructive? Can they lead to a system that
may convincingly develop a constructivist
system that becomes *cognitive*? I doubt it.

Natural environments

« 8 » What can be done about it? I be-
lieve that the constructivist community
should focus on the question of how sym-
bol processing capabilities can develop in
noisy, continuous environments – where
experiences are grounded and embodied in
an actual bodily perception-action system.
Evidence has been accumulating over re-
cent years that this is not an insurmountable
endeavor. The *theory of event coding* (Hom-
mel et al. 2001) postulates that events may
be a highly important cognitive concept for
structuring experiences and thus for per-
ceiving the environment in chunks that may
be symbolizable. Also, in the cognitive ro-
botics literature, the registration of events –
such as when touching an object – has been
acknowledged as one key mechanism for
segmenting the environment into meaning-
ful interaction components (Wörgötter et al.
2013). Bodily interactions with the environ-
ment were structured into a natural action
grammar with properties that are strongly
related to Noam Chomsky's universal gram-

column B

column C

mar (Pastra & Aloimonos 2012). Research
from my own group suggests that goal-ori-
ented representations should be separated
from representations of spatial interaction
for setting the stage to develop composi-
tional concept structures, which are neces-
sary for language development (Butz 2013).

« 9 » In conclusion, I agree with the
authors that schema learning approaches
should be re-considered and revived. Start-
ing with a symbolic world and facing one
particular, partially-observable toy problem,
however, will not advance schema learn-
ing mechanisms. Rather, these mechanisms
need to be implemented in environments
within which interactions are continuous,
state transitions are stochastic to a certain
degree, and perceptions are noisy. Tools and
mechanisms are currently being developed
that can segregate these continuous realms
into meaningful and purposeful symbol
systems. Key components of such mecha-
nisms are anticipations, modularizations,
and event-based separations. Measures of
valence and resulting purposeful, goal-ori-
ented interactions are most likely additional
key concepts. A learning system that builds
schemas based on these principles may in-
deed be the way forward towards scalable
cognitive systems that develop in complex
environments, effectively implementing
constructivist theories of cognition.

Martin V. Butz is Full Professor at the Eberhard Karls
University of Tübingen, in the Faculty of Science,
Department of Computer Science, and Department
of Psychology. His research group works on cognitive
modeling, focusing on how the brain develops
representations of the body and the surrounding
space and how these representations are used to
manipulate the environment goal-directedly.

RECEIVED: 16 OCTOBER 2013

ACCEPTED: 21 OCTOBER 2013

column C

column A

1 The Power of Constructivist 2 Ideas in Artificial Intelligence

3
4 Kristinn R. Thórisson

5 Reykjavik University, Iceland

6 thorrisson/at/gmail.com

7
8

9 > **Upshot** • Mainstream AI research
10 largely addresses cognitive features
11 as separate and unconnected. Instead
12 of addressing cognitive growth in this
13 same way – modeling it simply as one
14 more such isolated feature and continu-
15 ing to uphold a wrong-headed divide-
16 and-conquer tradition – a constructiv-
17 ist approach should help unify many
18 key phenomena such as anticipation,
19 self-modeling, life-long learning, and
20 recursive self-improvement. Since this is
21 likely to result in complex systems with
22 unanticipated properties, all cognitive
23 architecture researchers should aim to
24 implement their ideas in full as running
25 systems to be verified by experiment.
26 Perotto's paper falls short on both these
27 points.

28
29 « 1 » Cognitive growth, self-inspection,
30 anticipation (prediction based on partial
31 observation), self-organization – what do
32 these have in common? They are all part
33 of a growing set of concepts from biology,
34 cognitive science, artificial intelligence, and
35 psychology that must be related to one an-
36 other if we are ever to produce a coherent
37 theory of intelligence, whether in machines,
38 animals, or humans. And if our aim is to
39 build working systems – if our stance is a
40 software engineering one with an end-goal
41 of building deployable systems that can op-
42 erate in real-world environments, whether
43 it be space probes, housecleaning robots,
44 deep-sea explorers, or stock-market invest-
45 ment programs – then our methodological
46 approach must embody principles that are
47 useful for steering our efforts when design-
48 ing, architecting, implementing, and testing
49 our systems.

50 « 2 » Filippo Perotto presents in his paper
51 a model of an anticipatory learning mecha-
52 nism, CALM, which is based on construc-
53 tivist principles. His high-level model of
54 agent-environment coupling, CAES, seems
55 a reasonable one. Both models are based on

column A

column B

the fundamental assumptions, which I agree
with, that: (a) to understand intelligent be-
havior we must include in our analysis the
context in which it operates; and (b) most
environments of any interest to intelligent
beings contain a mixture of deterministic
and non-deterministic causal connections,
with many of the former remaining invis-
ible. In my view, and it would seem Perotto's
as well, an environment with complex causal
relationships (e.g., our everyday world)
gives rise to a vast number of potentially ob-
servable phenomena, many of which do not
clearly or readily convey their underlying
causes; this set of potential observable and
inspectable phenomena is nevertheless the
only information that an intelligent system
has access to, via their sensory apparatuses,
for anticipating how their external environ-
ment behaves so as to efficiently and effec-
tively achieve its goals within it.

« 3 » Before continuing with direct
commentary, some points are in order so
as to elucidate the context in which I look
at systems engineering, architecture, and
constructivism. Due to the high number of
combinatorics that a complex environment
will produce, through countless interactions
between its numerous elements, an agent
must create *models* that isolate and capture
some essence of underlying causes (invari-
ants or partial invariants) in this environ-
ment (Conant & Ashby 1970). Such mod-
els must be capable of capturing abstract
levels of detail that can be used to steer the
operations of a system towards efficient ex-
penditure of computational resources – any
thought spent on details completely unre-
lated to goals (future and present) would be
a waste of the agent's time. Thus, the partial
models of the environment that an intel-
ligent agent creates will likely form some
sort of a cognitive “random-access” abstrac-
tion hierarchy. Depending on the type of
current goal and situation, the agent can
then choose models at a particular level of
abstraction at any time to help it exclude
irrelevant issues from consideration when
decisions are being made about how to
achieve the goal in that situation. A coher-
ent, unifying model of cognition following
constructivist principles must explain how
this works, in particular how goals, models,
experiences, and iterative knowledge acqui-
sition and improvement operate in concert

column B

column C

to achieve cognitive growth in an agent. An
engineering methodology for how to build
artificial systems implementing such func-
tions must go further, by helping with de-
fining specifications for an implementable
architecture, and providing guidelines on
how to implement them in a way that allows
experimental evaluation.

« 4 » An artificial system built to
achieve general intelligence must be able to
deal with novel situations – situations not
foreseen by its programmers. Instead of be-
ing given pre-programmed algorithms by its
designers, known to be applicable to partic-
ular and specific problems, tasks, situations,
or environments, the AI itself must be im-
bued with the ability to generate algorithms
(or, compute a control function – I do not
distinguish between the two here). For this
to be possible, the system must further-
more be equipped with the ability to (re-)
program itself, otherwise it cannot sensibly
change its own operation in any meaningful
way based on acquired experience. And to
be able to do so, the system must be reflec-
tive – that is, the system's architecture and
operational semantics must be captured in a
way that enables it to read and interpret its
own structure and operation. This is what I
consider the essence of a constructivist AI
methodology: specifications for how to im-
bue machines with the capability to make
informed changes (whether slowly or quick-
ly) to their own operation, via the runtime
principles embodied in their architecture. I
do not believe constructivist AI can be done
without some form of self-programming on
the part of the machine, which in turn can-
not be achieved without transparency of its
operational semantics. In fact, even more
radically, I suspect artificial general intel-
ligence cannot be achieved *at all* without
such capabilities; higher levels of cognitive
operation in the context of novel or unan-
ticipated tasks, situations, and environments
must require some sort of cognitive growth
– namely, some form of re-programming of
the cognitive system's operation. Conversely,
constructivist views on cognition are so dif-
ferent and incompatible with standard soft-
ware engineering methodologies, especially
with its tradition of manual software crea-
tion, that they cannot be used at all for engi-
neering such systems. To address construc-
tivist principles head on in a computational

column C

column A

1 framework will require a new *constructivist*
2 *AI methodology* (CAIM; Thórisson 2012).

3 « 5 » Whether or not Perotto agrees
4 with my views on the nature and need for
5 constructivist development principles thus
6 outlined, he does make some claims to tak-
7 ing steps toward computational implemen-
8 tations of constructivist principles. In this
9 context, many important questions come
10 to mind – chief among them being how
11 effective the ideas are for explaining cogni-
12 tive growth in nature, and how useful might
13 they be for helping implement artificial gen-
14 eral intelligence. As Perotto's paper seems to
15 be aimed more at the second topic, we can
16 ask, firstly, do the ideas presented in his pa-
17 per help with – or are they likely to lead us
18 to – better software engineering methods
19 for implementing constructivist learning in
20 deployed systems? Secondly, we can ask, if
21 they do in fact offer some new insights to
22 this end, how much still remains to be ex-
23 plained for such systems to spring forth as a
24 result – or conversely, how big a part of the
25 constructivist puzzle does the work attempt
26 to address? Let us look at these in order.

27 « 6 » The aim of AI is not just to specu-
28 late but to build working, implemented sys-
29 tems. In AI, any theoretical construct aimed
30 at advancing our understanding of how to
31 implement cognitive functions should ulti-
32 mately be judged on whether actual imple-
33 mentation can conclusively, or partially, al-
34 low us to conclude through reliable means
35 (i.e., scientific experimentation), that the
36 ideas, when operating in a relatively com-
37 plete AI architecture situated in a complex
38 world (Perotto's target environments), are
39 capable of *scaling up*. By "scaling up" I mean
40 the ability of a system to grow in a way that
41 supports recursive self-improvement in
42 complex environments (e.g., the physical
43 world), with respect to its top-level goals.
44 This question is of course difficult to answer,
45 whether experimentally or analytically. A
46 quick walk down memory lane reminds us,
47 however, that the history of AI is replete
48 with examples of proposals that looked great
49 on paper but completely failed such scaling
50 up when implemented in a running system,
51 or when attempts were made to expand the
52 models the ideas embodied to include more
53 of the many functional characteristics that
54 they originally left untouched. Unfortunate-
55 ly, experimental evaluation of Perotto's pro-

column A

column B

posed ideas is touched on only briefly in the
paper, and the support provided to answer
this question is inconclusive at best. On this
count, therefore, not much can be said about
the scalability of Perotto's ideas. This is dis-
appointing because a fundamental feature
of known constructivist systems in nature is
their capability to grow cognitively with ex-
perience – itself a form of scaling-up. Other
phenomena, such as the power of the CALM
schema formalism to produce new knowl-
edge of complex environments, to support
models of self (required for any system ca-
pable of self-directed cognitive growth), and
their ability to support self-inspection, are
also not addressed to any sufficient extent
in the work. Since these issues are briefly
touched on or left unmentioned, we can
only assume that they remain unaccounted
for by the present work.

« 7 » My second question regards the
"size of the intelligence puzzle" addressed.
An artificial cognitive system must, to have a
chance at becoming a comprehensive theory
of the major facets of intelligence, include
a large number of functions that allow the
system to operate relatively autonomously in
complex environments. This *theoretical scal-*
ability of an isolated mechanism is its perse-
verance and robustness when included in a
better (larger, more comprehensive) model/
theory, which can in turn serve as the foun-
dation for building systems with increased
operating power, including an increased
capacity for cognitive growth and architec-
tural complexity. If Perotto's work turns out
to be correct, if it indeed offers, as Perotto
claims in the abstract, "a step towards com-
putational implementation of constructivist
principles," how much of the phenomenon
in question – cognitive growth – remains to
be explained? The lack of a clear connection
between his CALM and CAES models is al-
ready a sign that some amount of work re-
mains to be done in this direction. My own
list of candidate principles and features (cf.
some already mentioned above) that should
be accounted for in any reasonable theory
of cognitive growth is, unfortunately, quite
a bit longer than that addressed in Perotto's
paper. Firstly, as described above, cognitive
growth requires some kind of autonomic,
recursive self-improvement. Although my
team has made some progress on this front
recently (Nivel & Thórisson 2013, Nivel et

column B

column C

al. 2013), research on the topic is still in its
1 infancy, with a host of unanswered practical
2 and theoretical questions. Such questions
3 include: What kind of representations¹ are
4 amenable to automatic self-programming
5 for cognitive growth (existing programming
6 languages and paradigms created for hu-
7 mans require human-level intelligence to be
8 used – which calls for the very phenomenon
9 we are striving to understand how to imple-
10 ment); how can the transparent operational
11 semantics needed for automatic program-
12 ming be achieved? Related to that are the
13 questions: How can a system's operational
14 semantics be measured; what kind of meta-
15 level control structures can be used to steer
16 cognitive growth; what kinds of control ar-
17 chitectures can serve as host architectures
18 for the proposed (or any other) constructiv-
19 ist principles? Questions regarding theoret-
20 ical scalability issues loom large.

21 « 8 » These are, of course, not simple
22 topics. Quite the contrary, they are deep
23 and challenging. But they are central to
24 constructivist approaches, developmental
25 robotics, and principles of cognitive growth,
26 and it is precisely for that reason that they
27 must not be left unaddressed, lest our efforts
28 become victims to the same oversimplifi-
29 cation and incorrect application of divide-
30 and-conquer methodology that has plagued
31 much of AI research in the past half century
32 (cf. Thórisson 2013). Unlike so many other
33 phenomena in AI, e.g., planning, vision, rea-
34 soning, and learning, that have been largely
35 addressed by calling them "computational"
36 and studying them in isolation through the
37 same strictly allonomic methodologies as
38 used for banking systems, word processors,
39 and Web page construction, a constructivist
40 methodology holds a promise – a potential –
41 to unify a host of complex cognitive mecha-
42 nisms, most of which have eluded scientific
43 explanation so far. A holistic stance is by far
44 the most likely to lead to an understanding of
45 the phenomenon of intelligence, and anyone
46 with a constructivist mindset has already
47 taken an important step in that direction.
48 But for this to pave the way towards a bet-
49 ter

50
51 1 | My use of "representations" implies a
52 larger scope than models, capturing virtually
53 anything that might be needed to be encoded in a
54 particular runtime medium for a running ("live")
55 intelligent system.

column C

column A

1 ter theory, a genuine attempt must be made
2 to weave as many key cognitive phenomena
3 into the account as possible, to attempt to
4 provide a unifying account. And for any
5 engineering effort to be taken seriously, the
6 requirement for experimental evaluations of
7 (physical and/or virtual) running software
8 systems cannot go ignored. Perotto's stance
9 on these pressing issues remains for the time
10 being largely unknown; we can only hope
11 that he addresses them in the future.

12
13 **Kristinn R. Thórisson** has been doing research in
14 artificial general intelligence and real-time interaction
15 for over two decades in academia and industry. His
16 AERA constructivist cognitive architecture is the
17 world's first system that can learn complex skills by
18 observation in largely underspecified circumstances. He
19 is a two-time recipient of the Kurzweil Award and has
20 a Ph.D. from Massachusetts Institute of Technology.

RECEIVED: 16 OCTOBER 2013

ACCEPTED: 24 OCTOBER 2013

26 Anticipatory? Yes. 27 Constructivist? Maybe

28 Georgi Stojanov

29 The American University of Paris,
30 France

31 gstojanov/at/aup.edu

32
33
34
35
36 **Upshot:** The CALM cognitive agent with
37 its learning mechanism, as presented by
38 the author, can be described as “trivially
39 constructivist.” Probably, at best, it can be
40 seen as a model of the empirical abstrac-
41 tion but not of the reflective abstraction.
42 The “intrinsic motivations” in the simu-
43 lated agent presented as “evaluative
44 signals” sent from the agent's “body” to
45 its “mind” can be seen as low-level physi-
46 ological drives. They cannot account for
47 far more sophisticated intrinsic motiva-
48 tions such as curiosity.

49
50 «1» In the opening sections and in §1,
51 Filipo Perotto sets up a formidable challenge
52 for himself by promising a step toward an
53 artificial general intelligence (AGI) that fol-
54 lows the constructivist approach of Piage-
55 tian flavor. There is also an explicit critique

column A

column B

of this constructivist AI, in which despite all
the promises made, there has been a “lack
of concrete results.” The critique is justified
if the expected concrete result was to build
an artifact that would exhibit the behavior
of a three-year-old infant. We are certainly
not there yet. On the other hand, the con-
structivist AI approach certainly made huge
theoretical advances by demonstrating the
inappropriateness of the traditional soft-
ware methodology to deal with the design
of self-constructive autonomous intelligent
agents (e.g., Thórisson 2012), or shifting the
research focus to issues neglected in tradi-
tional AI: sensorimotor interaction, intrin-
sic motivation, complete cognitive archi-
tectures (e.g., Stojanov, Kulakov & Clauzel
2006; Stojanov & Kulakov 2011).

«2» In §2 and §3, Perotto introduc-
es the conceptual structure of a schema:
context \wedge action \rightarrow expectation, which he
also calls an “elementary piece of knowl-
edge.” The “context” vector represents the
readings of all external and internal sensors,
and when some “action” is executed, the
agent anticipates the outcome in terms of
the “expectation” vector. Thus, throughout
its lifetime, the agent put in particular en-
vironment should learn to predict the out-
comes of its actions (“to adapt itself”), even
if the environment is partially observable.
Many researchers have used this “context \wedge
action \rightarrow expectation” construct (Drescher
1991; Schachner 1996; Schachner, Real del
Sarte & León 1999; Tani 1996; Stojanov,
Bozinovski & Trajkovski 1997; Chaput 2004;
see Stojanov 2009 for an overview of com-
putational models of Piagetian schemas) in
the task of learning forward-models (or an-
ticipative models) of the environment. The
simulated environment is represented via a
FPOMDP (§28). The states of the environ-
ment are represented with a set of proper-
ties X, and among those properties there
are some that cannot be perceived by the
agent's perceptual apparatus. This leads to
perceptual aliasing and makes the problem
of learning effects of actions in given con-
texts much more difficult. CALM (§38) is
the learning mechanism designed to learn
the dynamics of the underlying FPOMDP
through execution of agent's actions (which,
from the point of view of the FPOMDP are
controllable variables) and construction
of reliable predictive schemas, described

column B

column C

above. §50, §51 and §52 describe the three
basic methods for schema construction in
CALM: differentiation, adjustment, and in-
tegration. As there are unobservable proper-
ties of the environment, sometimes CALM
will fail to predict accurately the effect of
some action, and in some cases, the situa-
tion can be remedied by abstract differen-
tiation (§55). Essentially, this means that
the context and expectation vectors are arbi-
trary values that are enlarged and attributed
in a way to make them distinct from existing
schemas. The new schemas are called syn-
thetic elements as they cannot be directly
perceived. The method of propagation of the
value of the synthetic elements is called ab-
stract anticipation (§57). Once (if the com-
plexity/observability ratio allows) the agent
using CALM learns the environment model
perfectly, it can always predict the effect of
its action in a given context.

«3» My condensed (and, I hope, not
too simplistic) description of CALM in
the previous paragraph is to show that (al-
though an original and efficient solution) it
is *constructivist* only in a trivial way: it learns
a model of its environment incrementally.
Perotto appears to be like many develop-
mental psychologists in the 1970s:

“What they [developmental psychologists]
called construction seemed to refer to the fact that
children acquire adult knowledge not all at once,
but in small pieces. I did not think that this was
a revelation and therefore called their approach
‘trivial constructivism.’” (Glaserfeld 2005: 10).

The monolithic single-thread algorithm is
completely deterministic and will (eventu-
ally) come up with the same result, given
the same learnable environment. There is no
learning-to-learn (i.e., change of the learn-
ing trajectory) nor ability for reconceptu-
alization of a given situation, or evolution
of more sophisticated intrinsic motivations
(more about motivations below). Moreover,
as Perotto notes in §9, “The agent needs to
be able to detect high-level regularities in
the dynamics of the environment, but this is
not possible if the agent is stuck in a rigid
representational vocabulary.” The represen-
tational vocabulary of a CALM-driven agent
is rigid: all of the possible different schemas.
Synthetic items definitely enlarge it, but only
up to a certain predefined limit. There is no

column C

column A

1 way in which (in a genuinely constructivist
2 spirit) the agent can impose some organiza-
3 tion on the sensed environment. In continu-
4 ation of §9 we can read

6 “In a constructivist approach, cognitive devel-
7 opment must be a process of gradual complexi-
8 fication of the intelligence, where more abstract
9 structures (symbolic) are built from simpler
10 sensorimotor interactions, in a way that harmo-
11 nizes the lived experiences with the constructed
12 model.”

13
14 CALM does not provide a way to build
15 “more abstract structures... from simpler
16 sensorimotor interactions.” At best, there
17 are the synthetic elements that contain ab-
18 stract properties in the sense that they do
19 not correspond to any sensory inputs. Given
20 that those abstract properties are added to
21 schemas whose context and action vectors
22 are equal, it is impossible to understand
23 them as abstract/symbolic structures in the
24 sense given in §9. In Piagetian parlance, the
25 learning exhibited by CALM could be seen
26 as model of the *empirical abstraction* but not
27 of the *reflective abstraction* that is crucial for
28 cognitive development and creative behav-
29 ior. Briefly, via *empirical abstraction*, some
30 quality (e.g., weight or color or contingency
31 among actions and qualities) is abstracted
32 from an object/situation. On the other hand,
33 reflective abstraction is about reorganiza-
34 tion of existing schemas and their projection
35 onto a higher plane. (See Kitchener 1986:
36 61–65 for an informative discussion of em-
37 pirical and reflective abstraction, as well as
38 the discussion in Campbell & Bickhard 1993
39 on the *knowing levels*).

40 « 4 » In §13, one can read that the
41 agent’s body with its “internal states and me-
42 tabolisms, elements that belong neither to
43 the mind nor the environment... allow the
44 agent to have intrinsic motivations...” I be-
45 lieve that the decision to introduce the two
46 entities (“body” and “mind”) is somewhat
47 arbitrary, given that it is barely mentioned in
48 the rest of the paper. It appears that the body
49 is introduced only to have the above-men-
50 tioned possibility to have “intrinsic motiva-
51 tions.” If this is the case, then the intrinsic
52 motivations can be related to low-level phys-
53 iological drives (hunger, pain-avoidance)
54 with no possibility for development of more
55 sophisticated forms of motivations such as

column A

column B

curiosity. If, on the other hand, the intrinsic
motivations can be placed in the “mind” of
the agent, I see no reason to draw the arbi-
trary body-mind distinction.

Georgi Stojanov obtained his PhD degree in
computer science and AI from UKIM University
in Skopje, Macedonia in 1997 on “Anticipation
theory and electro-pectograms in the context of
biological and machine intelligence.” His research
interests include: interactivism and constructivism;
cognitive development and creativity; analogy
and metaphor. He was a visiting scholar at the
University of Trieste, Italy, Les Archives Jean Piaget in
Geneva, Université de Versailles-Saint-Quentin-en-
Yvelines in Paris, and at the Institute for Non-linear
Science, UCSD. As of 2005, he is an associate
professor at the American University of Paris.

RECEIVED: 13 OCTOBER 2013

ACCEPTED: 18 OCTOBER 2013

Action, Anticipation, and Construction: The Cognitive Core

Mark H. Bickhard
Lehigh University, USA
mhbo/at/lehigh.edu

> **Upshot** • Interaction-based models of
cognition force anticipatory and con-
structivist models. The CALM model of-
fers significant development of such
models within a machine learning
framework. It is suggested that moving
to an entirely interactive-based model
offers still further advantages.

« 1 » Charles Sanders Peirce introduced
action and interaction as the proper loci for
understanding the mind well over a cen-
tury ago (Joas 1993). An interaction-based
model of cognition, in turn, is intrinsically
anticipatory – i.e., anticipations of potential
actions and interactions (Bickhard 2009b;
Buisson 2004; Piaget 1954). And an action
and interaction-based model of cognition
forces a constructivism: it is not feasible for
the world to impress competent interactive
system organization into a passive mind; it
must be constructed. For yet another step,

column B

column C

given that prescience does not exist, such 1
a constructivism must be a variation and 2
selection constructivism, an evolutionary 3
epistemology (Campbell 1974). These char- 4
acteristics, thus, form a coherent framework 5
for understanding cognition, and, more 6
generally, mind (Bickhard 2009b). 7

« 2 » Classic passive mind models, how- 8
ever, descending from the ancient Greeks, 9
still dominate the scene, currently in their 10
“recent” incarnations of symbolic compu- 11
tationalism and connectionism. Machine 12
learning is an interesting combination of 13
perspectives: learning about the environ- 14
ment requires checking what is tentatively 15
learned against that environment, which 16
requires action and anticipation and con- 17
struction of what is checked. Most cleanly, 18
what is checked are those anticipations per 19
se. But there is still also a reliance on passive 20
models of perception (generally based on 21
sensations) and restricted models of action 22
and construction. 23

« 3 » Filippo Perotto’s CALM is a signifi- 24
cant advance within this framework, espe- 25
cially in its ability to extract anticipatory in- 26
formation from an only partially observable 27
and not fully deterministic world, and to use 28
synthetic elements in doing so. It is impor- 29
tant to demonstrate that these more realis- 30
tic framework assumptions can be handled, 31
and to show how they can be handled. 32

« 4 » But CALM, too, is built on sensa- 33
tion models of perceiving and on singleton 34
actions. One of the current foci for devel- 35
opment of the CALM model is to develop 36
possibilities of chaining schemas – again, 37 317
I would agree that this is exactly the right 38
direction. I would like to comment, how- 39
ever, on an even more general approach that 40
might be considered – a fully interactive ap- 41
proach. 42

« 5 » Consider that passive sensations, 43
insofar as they exist at all, functionally serve 44
to detect properties of the environment, and 45
that such detection – as a strictly factual 46
matter – is all that is functionally relevant 47
to the system. In particular, such detections 48
need not be understood to *represent* that 49
which is detected in order to account for 50
their influences on system processing. Still 51
further, such detections can also be real- 52
ized by fully interactive processes, not just 53
by passive receptive processes (Bickhard & 54
Richie 1983). On the other hand, anticipa- 55

column C

column A

1 tions concerning possible interactions with
2 the environment also can, and arguably do,
3 occur with respect to *whole patterns* of inter-
4 action, not just singular actions. Chaining of
5 schemas is precisely a step in this direction,
6 but it requires more than “chains” to be able
7 to model general interactive patterns.

8 « 6 » So, I would suggest that patterns of
9 interaction can serve:

10 1 | detection functions, rather than sensa-
11 tions and perceptions interpreted as
12 representational (with all of the classic
13 problems that that interpretation entails:
14 Bickhard 2009b),² and

15 2 | as patterns of interaction that are antici-
16 pated as possible in the future, and

17 3 | as patterns that can be tentatively con-
18 structed in learning more about the envi-
19 ronment – learning more about what
20 patterns of interaction can be anticipat-
21 ed as possible, given what prior interac-
22 tive detections have already occurred.

23 « 7 » Such shifts generate a dynamic
24 systems model, more than a classic compu-
25 tational model, but one in which represen-
26 tation is not absent: truth value emerges in
27 anticipations that are capable of being true
28 or false, and cognitive representation more
29 generally can be built from organizations of
30 such anticipations (Piaget 1954; Bickhard
31 2009b). In such a model, representation is
32 not built on or out of presumed sensations
33 as representations.

34 « 8 » Modeling the dynamics of such
35 dynamic systems is difficult. For one class
36 of problems, there are no topological or
37 metric spaces built in to serve as spaces for
38 generalization. On the other hand, if the
39 construction of such topological spaces can
40 itself be constructed, then we can model the
41 cognitive development and organization
42 and re-organization of such spaces in chil-
43 dren and adults – a higher level of learning
44 and development than is usually addressed
45 (Bickhard & Campbell 1996). For another
46 class of problems, cognitive representations
47 of, for example, objects or numbers, cannot

48
49 2| Note that this also frees the model from
50 being able to generalize only along the dimen-
51 sional variables that are built into the sensation
52 apparatus, and from the related built-in metric
53 spaces for error, etc. Of course, it also makes the
54 dynamics of such generalization more difficult to
55 model.

column A

column B

be simply presupposed, but must themselves
be constructed. However, that was among
the basic insights of Piaget some time ago
(Piaget 1954; Allen & Bickhard 2011, 2013a,
2013b, 2013c).

« 9 » Overall, then, moving to a fully
interactive dynamic systems framework
makes a number of modeling problems
much more difficult. But it offers advantages
of avoiding classic problems concerning, for
example, the nature of representation (Bick-
hard 2009b), and offers direct approaches to
modeling phenomena that are very difficult
to approach within standard frameworks
(e.g., re-organizing the topology of repre-
sentational spaces in response to under-
standing an analogy; Bickhard & Campbell
1996).

Mark Bickhard is the Henry R. Luce Professor in
Cognitive Robotics and the Philosophy of Knowledge at
Lehigh University. He is affiliated with the Departments
of Philosophy and Psychology, and is Director of
the Institute for Interactivist Studies. His work
focuses on the nature and development of persons,
as biological, psychological, and social beings.

RECEIVED: 11 OCTOBER 2013

ACCEPTED: 21 OCTOBER 2013

Representing Knowledge in a Computational Constructivist Agent

Thomas Degris

ENSTA ParisTech-INRIA, France

thomas.degris/at/inria.fr

> **Upshot** • The aim of this commentary
is to relate the target article to recent
work about how to represent the knowl-
edge acquired from experience by a con-
structivist agent.

« 1 » Constructivist agents acquire new
knowledge and maintain existing knowl-
edge by experimenting with their environ-
ment. A key question is then how to repre-
sent knowledge for such an agent.³ In the

3| Joseph Modayil and I proposed an answer
to that question in their presentation “Scaling-up

column B

column C

target paper, knowledge that can be created
and updated from data is emphasized, but
a different mathematical framework and a
different architecture, namely the Horde ar-
chitecture (Sutton et al. 2011), is used. This
commentary presents the similarities and
differences between the target paper and the
Horde architecture.

« 2 » Both papers focus on a situated
agent embedded in its environment. The
agent does not have access to the full state of
the environment. To be able to understand
better its interaction with the environment,
the agent needs to construct abstract inter-
nal structures from a low level sensorimotor
loop. In both papers, the internal represen-
tation built by the agent comes from its own
experience and does not need to match an
arbitrary absolute representation of its envi-
ronment.

« 3 » The target paper has chosen to rep-
resent the knowledge of the agent with tree-
structured representations. While trees can,
in principle, take advantage of specific struc-
tures in the data, they also have issues that
can make them impractical to use as a life-
long constructivist agent in the actual world.
More specifically, as mentioned in the target
paper, the main idea of structured represen-
tations is that the system dynamics can be
factored to save memory and computational
time. But such structure just may not be in
the data. For instance, if one would like to
predict the next value of a bump sensor on
a small mobile robot, it is likely that infor-
mation from all the sensors on the robot, as
well as many abstract representations, may
help in one way or another to make a better
prediction. Thus, a prediction as simple as
the value of a bump sensor may simply not
be factorable. Moreover, even when some of
the system dynamic may be factorable, there
is no guarantee that other representations,
such as value functions or policies, will be
factorable (Boutillier, Dearden & Goldszmidt
2000). For an agent to take complex decision
or to understand a complex environment,
perhaps it is unavoidable to consider a large
number of variables or signals. In compari-
son, the Horde architecture focuses on al-
gorithms with a complexity in computation

knowledge for a cognizant robot” at the AAAI
Spring Symposium on Designing Intelligent Ro-
bots: Reintegrating AI, Stanford University, 2012.

column C

column A

1 and memory that is linear to the number of
2 parameters to learn. In practice, thousands
3 of predictions depending on thousands of
4 features can be learned online in real time
5 on an actual robot (Modayil, White & Sut-
6 ton 2012).

7 « 4 » The agent in the target paper con-
8 structs a set of schemas to build a predictive
9 representation. A schema takes a context
10 and an action to make a prediction about
11 the next time step. A context can be seen as
12 a set of conditions on the agent state; that is,
13 conditions on internal variables and the last
14 observation from the sensors. The action in
15 the schema describes what the agent will do
16 to go to the next time step. Thus, knowledge
17 constructed by the agent answers questions
18 such as: “Am I going to be connected to my
19 docking station at the next time step if I do
20 this action?” In comparison, demons in the
21 Horde architecture can represent knowl-
22 edge similar to schemas but also more gen-
23 eral knowledge: for a given agent state and a
24 policy – that is, a sequence of (stochastically
25 chosen) actions – a demon makes a tempo-
26 rally abstract prediction. For instance, an
27 agent can construct the answer to questions
28 such as: “What is the probability of being
29 connected if I follow the policy going back to
30 my docking station?” or “How much energy
31 will I use to go back to my docking station?”
32 Moreover, there are two additional advan-
33 tages with temporal abstractions. First, de-
34 mons can be used to build predictive fea-
35 tures in the agent internal state. Predictive
36 state representations (PSRs) are known to
37 be more general than POMDPs or n^{th} -order
38 Markov models – representations based on
39 history (Singh, James & Rudary 2004). Sec-
40 ond, it becomes possible to consider high-
41 level planning on temporal abstractions, as
42 has been proposed with the option frame-
43 work (Sutton, Precup & Singh 1999).

44 « 5 » Overall, the Horde architecture
45 has two key features compared to the rep-
46 resentation used in the target paper. First,
47 Horde can learn and maintain predictive
48 knowledge online and in real time. Second,
49 Horde can learn answers to temporally-ab-
50 stract questions. Thus, the Horde architec-
51 ture is a direct possible answer to two of the
52 questions mentioned at the end of the target
53 paper: how to extend the work to stochas-
54 tic and continuous environments and how
55 to consider action sequences. Of course, the

column A

column B

Horde architecture asks its own set of ques-
tions: What are the criteria to create or de-
lete demons based on data and experience?
What should be the behavior of an agent in
its environment to optimize learning in de-
mons (intrinsic motivations)? The path to a
constructivist agent for a general artificial
intelligence remains uncertain.

Thomas Degris is a postdoctoral fellow in the Flowers
team from ENSTA ParisTech-INRIA. He did his PhD at
University Pierre et Marie Curie on factored Markov
decision processes. He also worked as a postdoctoral
fellow in the RLAI group at the University of Alberta.

RECEIVED: 13 OCTOBER 2013

ACCEPTED: 18 OCTOBER 2013

Some Comments on the Relationship Between Artificial Intelligence and Human Cognition

Bernard Scott

Centre for Sociocybernetics
Research, Bonn, Germany
bernces1/at/gmail.com

> **Upshot** • In making a contribution to
artificial intelligence research, Perotto
has taken note of work on human cogni-
tion. However, there are certain aspects
of human cognition that are not taken
into account by the author’s model and
that, generally, are overlooked or ignored
by the artificial intelligence research
community at large.

« 1 » In his paper, Filippo Perotto has
taken note of work on human cognition.
In particular, he references Jean Piaget (§6)
and Ernst von Glasersfeld (§7). The former
developed his “genetic epistemology” by
studying the development of human chil-
dren. The latter, using Piaget as one of his
main sources, developed “radical construc-
tivism,” a philosophical treatise about how
humans come to know. Rather than at-
tempt to position the author’s work in the
broad field of artificial intelligence research,
something I do not feel confident to do

column B

column C

without further reading, I wish to note as-
pects of human cognition that the author’s
model does not take into account and that,
generally, are overlooked or ignored by the
artificial intelligence research community at
large.

« 2 » First, I note that humans, like
other biological organisms, are dynamical
systems (§11), far from equilibrium, whose
structures are continually being formed and
reformed by the dissipation of energy. The
author does state that humans are dynamical
systems; however, his account is limited
to the statement (in footnote 3) that “A dy-
namical system consists of an abstract state
space evolving in time according to a rule
that specifies the immediate future state
given the current state.” Far richer concepts
of what dynamical systems are and the chal-
lenges of modeling them are to be found, for
example, in the writings of Heinz von Fo-
erster (2003: chapter 1) and Ilya Prigogine
(1981) on self-organisation. Humans are
also organisationally closed, autopoietic
systems, endowed with an operationally
closed nervous system. Using these foun-
dational ideas, Humberto Maturana and
Francisco Varela (1980) developed a “biol-
ogy of cognition.” This work adds consid-
erably to our understanding of constructive
cognitive processes. Related ideas are to be
found in chapters 8, 10 and 11 of von Fo-
erster (2003), where there is discussion of
how sensorimotor activity leads to the com-
putation of “objects” as an invariant of an
organism’s constructed reality. These works
are seminal accounts of what is referred to
in later literature as “enactive cognition.”

« 3 » The author uses the term “sym-
bolic” in §9. The author’s model is presented
as a general mechanism for learning. It, like
much other work in artificial intelligence
research, ignores or takes for granted that
which is, with few exceptions, peculiarly
human in human cognition: the ability to
communicate and compute using what
George Herbert Mead refers to as “signifi-
cant symbols” – gestures, icons and utter-
ances that call forth in the sender similar
response to those elicited in the receiver.
Humans converse with each other and con-
verse with themselves. This truth falsifies
the claim made by many in the artificial in-
telligence community that brains and com-
puters are both “physical symbol systems.”

column C

column A

1 Other critics (e.g., Searle 1980) have also
2 challenged this claim. Scott & Shurville
3 (2011) provide an extended discussion of
4 the topic and propose its falsification based
5 on their analysis that a “symbol” is a second-
6 order “object” that two or more interacting
7 organizationally closed systems compute as
8 standing for a given first-order “object” and
9 compute that they are both doing so.

10 « 4 » In his model, the author of the
11 target paper refers to his simulated agent as

column B

having a “mind” (§13). If we take “mind” to
refer to the conceptual processes that con-
stitute humans as individual selves, then it
is possible find in the literature more elabo-
rated understandings of “mind” as an em-
bodied, organizationally closed, self-repro-
ducing system of concepts that arises as a
consequence not only of ongoing cognitive
constructions but also of social interaction
(Pask, Scott & Kallikourdis 1975; Pask 1981;
Scott 2007).

column C

Bernard Scott completed a Ph.D. in cybernetics from
Brunel University, UK, in 1976. His supervisor was
Gordon Pask, with whom he worked between 1967
and 1978. Among other positions, Bernard is a Fellow
of the UK’s Cybernetics Society and Past President
of Research Committee 51 (on sociocybernetics)
of the International Sociological Association.

RECEIVED: 14 OCTOBER 2013 9

ACCEPTED: 31 OCTOBER 2013 10

Author’s Response: Evaluating CALM

Filipo Studzinski Perotto

> **Upshot** • In this response, I address the
24 points raised in the commentaries, in
25 particular those related to the scalability
26 and robustness of the mechanism CALM,
27 to its relation with the CAES architecture,
28 and to the transition from sensorimotor
29 to symbolic.

General claims

32 « 1 » The commentators have touched
33 important points in the ideas presented
34 in the article. Some of the criticisms made
35 might appear heavy, but this is due to the na-
36 ture of this research, not limited to technical
37 applied AI questions, which aims to address
38 challenging philosophical and scientific
39 problems.

40 « 2 » Since I declared in the beginning
41 of the paper that until now, Constructivist
42 AI has not been able to present “impressive
43 results,” I led the readers to expect some
44 spectacular results. However, the stated ex-
45 perimental outcomes (with the hyper-flip
46 problem) can rather disappoint such ex-
47 pectations. The assertion might also give
48 the wrong impression that I considered
49 Constructivist AI stagnant until the arrival
50 of this article. I am in complete agreement
51 with **Georgi Stojanov** when he says that con-
52 structivist AI “certainly made huge theo-
53 retical advances” (§1), and I would add that
54 AGI has incorporated several concepts from
55 the constructivist approach, even if those

column A

researchers do not necessarily call them-
selves constructivists. **Stojanov** says that such
critique would be justified “if the expected
concrete result was to build an artifact that
would exhibit the behavior of a three-year-
old infant” (§1). To date, we are not able to
do so, neither within the constructivist ap-
proach, nor with any other form of AI.

« 3 » In the long road towards con-
structivist artificial general intelligence, my
article aims to be just a step forward, but it
is still far from the finishing line. The ideas
presented make up just a further brick for
constructing the bridge, and not a complete
definitive answer. As **Stojanov** says, this is al-
ready “a formidable challenge.”

« 4 » As is often the case with most of
these ambitious investigations, the work
done until now left more open hypotheses,
unanswered questions, ideas and promises,
than actually determined conclusions or
remarkable results. Nevertheless, thanks to
that ambition, it is possible to believe that
the work done, albeit quite modest, points
in the right direction.

« 5 » Despite all the efforts, we still
find ourselves stuck between two steep
challenges. On the one hand, there is the
complexity of the sensorimotor problems,
which require computationally viable mod-
els capable of treating large continuous
domains and realizing cybernetic adapta-
tion, interactive processing of imprecision,
refinement of skills, etc. On the other hand,
there is the problem of constructing sym-
bols to represent abstract entities and pro-
cesses, which could lead the agent to a kind
of higher level of thought in which the ex-

column B

perience is organized in terms of intelligible
concepts. The ideas presented in my article
do not solve either of these challenges but
could eventually help AI to get a foothold
in both.

Scalability and robustness

« 6 » The first important question cited
many times in the commentaries can be
summarized like this: can the mechanism
scale up well to complex, continuous, large-
order, real-time, noisy, non-deterministic
environments? In other words: can the vi-
ability of the model be convincingly demon-
strated in an experimental way? As claimed
in the introduction to my article, so far no-
body has been able to do this in constructiv-
ist general artificial intelligence.

« 7 » CALM, too, suffers from scalabil-
ity difficulties. It can scale up well on highly
structured environments, where the agent
deals with a large number of variables but
where causal links are very precise, where
relevant variables in function of the agent’s
goals are easy to identify, and where non-ob-
servable variables exist on a very small scale.
I agree that it is easy to be robust in such
environments. Since CALM was designed
to work in discrete symbolic environments,
it is not adapted to be directly applicable to
large sensorimotor problems.

« 8 » **Frank Guerin** (§6) suggests the ex-
ample where two stereo cameras deliver a
few million pixels in 24 bit color at thirty
frames per second and CALM tries to pre-
dict the consequences of actions in the com-
plexity of an everyday setting. He wonders
whether each bit of input could be used as

column C

column A

1 a CALM variable. Admittedly, CALM is not
2 adapted to face this kind of problem.

3 « 9 » I believe that for CALM to deal
4 with sensorimotor problems of larger mag-
5 nitudes, a bigger architecture must be devel-
6 oped. This includes tools that can segregate
7 “continuous realms into meaningful and
8 purposeful symbol systems” (Martin Butz
9 §9) processing sensorimotor signals before
10 linking them with other CALM modules. In
11 humans, what is conveyed from the eyes to
12 the association areas is more than a matrix
13 of pixels, and is rather information about
14 lines, contours, contrasts, movements, basic
15 forms, etc.

16 « 10 » The human brain is not a flat
17 system processing all signals at once, but is
18 divided in several zones and layers that are
19 more or less specialized. Sensory and motor
20 data are processed in primary areas before
21 being integrated in the association zones
22 of the neural cortex (Tortora & Derrickson
23 2012). Roughly speaking, perhaps CALM
24 can be related to the association cortex rath-
25 er than to the sensorimotor cortex. Guerin
26 puts the right question when asking “where
27 to make the cut-off between what the core
28 CALM system sees and what is the respon-
29 sibility of other abstraction mechanisms”
30 (§6). For now, this question must remain
31 unresolved.

Experimental scenario

32
33
34 « 11 » The experimental problem used in
35 my article (hyper-flip), although adequate
36 for illustrating the mechanism's capabili-
37 ties, is admittedly too simplistic. It remains a
38 tricky toy problem, which can demonstrate
39 neither the algorithm's robustness nor how
40 to solve concrete problems. It would be
41 necessary to conduct more sophisticated
42 experiments to show that CALM could be
43 able to discover unobservable and relevant
44 environmental properties, representing and
45 using them efficiently as synthetic elements
46 in its world model when facing more com-
47 plex problems.

48 « 12 » I can only agree with Kristinn
49 Thórisson when he says that “the aim of AI
50 is not just to speculate but to build working,
51 implemented systems” (§6) and that “for any
52 engineering effort to be taken seriously, the
53 requirement for experimental evaluations of
54 (physical and/or virtual) running software
55 systems cannot go ignored” (§8).

column A

column B

« 13 » I am also in agreement with Butz
when he declares in the upshot that “sensory
as well as motor noise is ubiquitous in our
environment” and that “symbols do not ex-
ist a priori but need to be grounded within
our continuous world.” Simple high-level
symbol manipulation problems that ignore
the low-level sensorimotor challenges can
hardly lead to a system that may convinc-
ingly become cognitive.

« 14 » Like several other researchers, I
believe that the domain *par excellence* for
testing machine learning models is robot-
ics. A simple robot with continuous noisy
sensors in real-time action into the physical
world is a fantastic challenge for such gen-
eral AI systems. In its current stage, CALM
is not yet ready to face that kind of problem
successfully. In the following paragraphs, I
would like to address some of the directions
I can envisage it taking in order to move for-
ward.

From continuous signals to discrete representations

« 15 » One of the main limitations of
CALM is the need for a predefined discrete
representation of both the signals received
and those transmitted by the agent. Also,
time is considered as a discrete succession
of cycles. However, many problems in com-
plex environments can only be properly
addressed through continuous representa-
tions, which enable an agent to face prob-
lems on the sensorimotor level.

« 16 » It seems more natural to start
with continuous signals and gradually con-
struct discrete states as a sort of abstraction.
This is the first step to going beyond sensori-
motor primitives and arriving in a symbolic
dimension. It also applies to temporal ab-
straction because intelligence needs to slice
the continuous flow into relevant pieces of
time in order to recognize events or cycles.

« 17 » In any case, schema learning
mechanisms are not necessarily incompat-
ible with continuous environments. An ex-
tension of the schema used by CALM can
be used to represent changes in continuous
variables. Very basically, we can represent
the anticipation of an increase or decrease in
the value of a certain variable, or the tenden-
cy to converge towards some specific value,
given some action. In this way, each schema
realizes a kind of simplified regression,

column B

column C

where in function of some conditions (con- 1
text and action), the schema can anticipate a 2
continuous variation of some variable. 3

Noise and non-determinism

« 18 » The definition and exploration 6
of environments that I called “partially de- 7
terministic” (e.g., §2) should be considered 8
worthwhile. The methods behind CALM 9
were defined to focus on the discovery of 10
deterministic regularities in an environment 11
composed of deterministic and non-deter- 12
ministic phenomena. 13

« 19 » For an agent, a complex environ- 14
ment can appear non-deterministic because 15
its perception, control and understanding 16
are limited in some way (partial observabil- 17
ity, noisy sensors, imprecise effectors, other 18
entities acting in the same environment, too 19
many complex causal relations, etc.). Appar- 20
ent non-determinism can be modeled either 21
by creating stochastic rules, or by continu- 22
ing to search for causes. 23

« 20 » Every roboticist knows the im- 24
portance of taking the noise and impreci- 25
sion inherent in sensory and motor appa- 26
ratus seriously. So far, CALM has not been 27
equipped with any mechanism to treat 28
noise explicitly. Even so, the presence or 29
absence of noise could be represented as a 30
cause of some perturbed anticipations, as in: 31
 $a \wedge \sim \text{noise} \rightarrow b$. 32

« 21 » That said, the possibility of repre- 33
senting certain situations as stochastic regu- 34
larities could be incorporated into CALM, 35
working as a complementary method for 36
situations where deterministic assumptions 37 321
are not possible. Such a method can allow 38
the mechanism to search for probabilities 39
in order to anticipate which is most likely to 40
happen. Nevertheless, that kind of search for 41
statistical regularities should not interfere 42
with the search for causal relations. 43

Robustness and parallelization

« 22 » Thomas Degris says that for “an 46
agent to take complex decisions or to under- 47
stand a complex environment, perhaps it is 48
unavoidable to consider a large number of 49
variables or signals” (§3). This is certainly 50
correct, especially when the problem is close 51
to the sensorimotor level. 52

« 23 » As Stojanov claims, CALM resem- 53
bles a completely deterministic “monolithic 54
single-thread algorithm” (§3). In nature, 55

column C

column A

1 animal as well as human brains do not oper-
2 ate as a centralized hierarchy, but more like
3 cooperative and concurrent modules work-
4 ing simultaneously in several levels, and not
5 necessarily in complete, harmonious coher-
6 ence. The power of intelligence stems from
7 the diversity of many effective imperfect
8 methods, and decisions emerge from con-
9 flicts and negotiations among them (Min-
10 sky1988: 308).

11 « 24 » I believe that any good construc-
12 tivist AI program will have to end up being
13 more or less in accordance with that system-
14 ic modular perspective of the mind, where
15 learning will appear as a continual construc-
16 tion and reconstruction of modules, each
17 one working in a specific level and domain,
18 but in constant interaction with other mod-
19 ules. Once within this conjuncture, CALM
20 could be imagined as the engine inside some
21 modules, under the baton of some principle
22 responsible for coordinating the modules in
23 the whole system.

24 « 25 » Moreover, concerning robust-
25 ness, parallelization is a very powerful
26 means to break complexity and to deal with
27 complex environments. The neural organi-
28 zation and functioning of the brain is highly
29 parallelized. Although it was not mentioned
30 in my article, CALM can implement a kind
31 of parallelization, since the construction of
32 each anticipatory tree (that models the dy-
33 namics of one single variable) can be real-
34 ized independently from the other trees, i.e.,
35 in different separated threads.

36 « 26 » Another way to be robust is to
37 pay attention to what is important (Foner
38 & Maes 1994). The problem of indistinctly
39 correlating actions with changes in sensor
40 data is computationally unfeasible for any
41 non-trivial application. This problem be-
42 comes more manageable by restricting the
43 set of sensor data the agent attends to, or the
44 set of internal structures that is updated, at
45 particular instants. In the same vein, CALM
46 implements a focus of attention related to
47 the affectively important variables.

48 « 27 » Degrís writes that “while trees
49 can, in principle, take advantage of specific
50 structures in the data, they also have issues
51 that can make them impractical to use as
52 a life-long constructivist agent in the ac-
53 tual world” (§3) and that “even when some
54 of the system dynamic may be factorable,
55 there is no guarantee that other represen-

column A

column B

tations, such as value functions or policies,
will be factorable” (ibid). Some technical
choices with regard to CALM’s methods
should be revised, especially with regard
to the management of episodic memory
and anticipatory trees. It is evident that a
robust algorithm for such general purposes
must be carefully studied. In other words,
the algorithms in CALM will probably need
certain improvements.

« 28 » Degrís cites the “Horde” archi-
tecture, suggesting that it can “represent
knowledge similar to schemas but also
more general knowledge” (§4). Horde can
construct “demons,” which are generalized
value-functions for given partial policies.
Those demons can be learned in parallel by
an efficient extended reinforcement learning
method during the actuation of the agent.
However, I think that the knowledge repre-
sented by Horde is not that similar to what is
represented by CALM.

« 29 » Space does not allow for a more
detailed comparison between CALM and
Horde. However, it is evident that architec-
tures like Horde will be a precious source of
good strategies for dealing with large real-
time sensorimotor problems, translating
them, when necessary, into symbolic terms.
I believe that, the crucial problem of using
efficient forms of representation aside, the
most important challenge is to find a way
to connect consistently the sensorimotor
(continuous, noisy, real-time, large scale)
domains to basic symbolic domains, and the
latter to more abstract ones.

From lower to higher levels

« 30 » Another major question repeat-
edly mentioned in the commentaries is the
passage from lower levels of interaction,
based on sensorimotor primitives, to higher
levels, based on abstract concepts. The ques-
tion can be formulated like this: Is CALM
able gradually to construct successive lay-
ers of abstraction in order to represent its
knowledge?

« 31 » According to Jean Piaget (1957),
from a fragmented sensorimotor universe,
intelligence builds elementary notions, de-
fines relations, finds regularities and even-
tually constructs an objective, substantial,
spatial, temporal, regular and external uni-
verse, independent of the subject itself. A
subjective “reality” will emerge from the in-

column B

column C

creasing coherence between schemas in the
course of these adaptations.

« 32 » In §3, Stojanov claims that CALM
does not provide a way to build more ab-
stract structures from simpler sensorimo-
tor interactions. At least he recognizes that
CALM is able to create synthetic elements
that enlarge the sensorial context with
something that is beyond perception. Even
if this is simple, the synthetic elements are
certainly a form of abstraction since they do
not correspond to any sensory input. How-
ever, once CALM places the synthetic ele-
ments side by side with the sensorial ones, it
does not create layers. The context is repre-
sented as a single flattened array. Evidently,
we cannot go too far without some kind of
robust structuring mechanism in order to
organize knowledge into different levels.

« 33 » In human beings, cognition is in
some way the construction of several lay-
ers of abstraction in order to understand
and interpret experiences. If this process is
compared with flying from the Earth to the
Moon, the inference of synthetic elements
would correspond to the takeoff. It does not
give us too much altitude but it is crucial to
start the voyage.

« 34 » Building synthetic elements does
not constitute a form of abstract or sym-
bolic thought by itself, but such a process
contains the basic insight of what we could
call “concept invention.” Synthetic elements
allow the designation of entities that cannot
be represented from combinations of direct
sensory perceptions. Thus, the possibility
of representing unobservable conditions is
a breakthrough along the road from mere
direct perception to more abstract forms of
understanding.

Grounding symbolic concepts on sensorimotor flows

« 35 » Guerin correctly claims that high-
level cognition is very much grounded in
sensorimotor intelligence (§4). I believe
that extracting significant symbolic con-
cepts from interactive sensorimotor flows is
one of the key challenges for AI today. The
robotics community and the symbolic AI
community can be seen as digging tunnels
on the opposite sides of a mountain. Despite
a lot of progress, a consistent integration of
contributions from the two sides is still in-
cipient. The same metaphor can be used to

column C

column A

1 refer to the relation between neuroscientists
2 and psychologists. As pointed out by **Guerin**,
3
4 “there are some works in cognitive science that
5 are beginning to attempt to address the issue of
6 providing some theoretical framework to account
7 for how a sensorimotor level can connect with
8 higher levels of cognition.” (§5)

9
10 In fact, the search for mechanisms capable
11 of doing so will draw on the findings from
12 all these fronts: high- and low-level, compu-
13 tational and cognitive sciences.

14 « 36 » Moreover, as **Thórisson** says,

15
16 “due to the high number of combinatorics that
17 a complex environment will produce, through
18 countless interactions between its numerous ele-
19 ments, an agent must create *models* that isolate and
20 capture some essence of underlying causes.” (§3)

21
22 Following constructivist principles, I would
23 suggest that the passage from the sensori-
24 motor to the conceptual domain is possible
25 through a series of abstractions where, at
26 each step, a large number of localized, con-
27 text-dependent, quick and small elements
28 are coordinated in more general, independ-
29 ent elements. Because CALM is not able to
30 do so, **Stojanov's** claim in his upshot is cor-
31 rect: it can be seen, at best, “as a model of the
32 empirical abstraction but not of the reflec-
33 tive abstraction.”

34 « 37 » In the same vein, **Mark Bickhard**
35 says that “CALM... is built on sensation
36 models of perceiving and on singleton ac-
37 tions.” (§4). I am in agreement with him
38 when he claims that “anticipations concern-
39 ing possible interactions with the environ-
40 ment... occur with respect to *whole patterns*
41 of interaction, not just singular actions”
42 (§5).

43 « 38 » Enabling the mind of an agent to
44 learn and think in terms of “whole patterns
45 of interactions” is another major challenge
46 in AI. These patterns must be related in two
47 ways:

48 1 | spatially, to high-level constructed ob-
49 jects, and
50 2 | temporally, to abstract events, i.e.,
51 “cognitive concepts for structuring ex-
52 periences and thus for perceiving the
53 environment in chunks that may be
54 symbolizable” (**Butz §8**) or “temporally
55 abstract prediction[s]” (**Degrís §4**).

column A

column B

CALM and CAES

« 39 » **Thórisson** points to the lack of
a clear connection between CALM and
CAES models. **Stojanov** claims that the de-
cision to introduce the two entities (body
and mind) was somewhat arbitrary. Fur-
thermore, **Bernard Scott** expressed his dis-
appointment with the way the term “mind”
was employed in my article. I agree that the
relation between CALM and CAES was not
developed in the paper, and that the hyper-
flip experiment does not illustrate that
relation. So let me try to make up for this
omission.

« 40 » CAES is an architecture that
connects concepts from cybernetics, the
theory of autopoiesis, dynamical systems,
and affective AI. It is based on the defini-
tion of three entities: environment, body,
and mind. CALM is the engine that plays
the role of the cognitive system in the
mind. Besides a cognitive system, the mind
includes an affective system (responsible
for evaluating the perceived situations), an
emotional system (directed to the internal
body states), and a reactive system (direct-
ed to the body effectors).

« 41 » Ross Ashby (1952) defined the
organism (or the agent) as a system com-
posed of a set of essential variables that
must stay within a certain physiological
normality (limits of viability) in order to
preserve the system's integrity and, conse-
quently, the organism's survival. A given
behavior contributes to the agent's adap-
tion if it ensures the persistency of these es-
sential variables within its viable limits. The
presence of essential variables assumes that
the agent has something like an internal en-
vironment. That is the body (Parisi 2004).

« 42 » In nature, organism and envi-
ronment can exert opposing forces with
respect to the global system's flow. How-
ever, only the organism is at risk of disin-
tegration, of disappearing as unity. A non-
destructive dynamical coupling is reached
in the relation between the two systems
when the organism interacts with the en-
vironment in order to ensure its self-pres-
ervation.

« 43 » Randall Beer (1995) integrated
the cybernetic concept of organism and
autopoiesis using dynamical systems. The
adaptation criterion is abstractly represented
as a zone in the space where the flow of the

column B

column C

system must remain. The limits of adaption 1
are the frontiers of that region within the 2
global system space (composed by agent 3
and environment), and the agent is con- 4
sidered adapted to the environment if its 5
activity drives the global system's trajectory 6
in such a way that it never escapes from 7
those frontiers. 8

« 44 » I agree with **Scott** when he says 9
that “humans, like other biological or- 10
ganisms, are dynamical systems, far from 11
equilibrium.” Even if my article does not 12
address this issue, CAES architecture was 13
imagined to correspond to a definition of 14
an agent as a system far from equilibrium, 15
in the sense proposed by Bickhard (2009a) 16
or Xabier Barandiaran and Alvaro Moreno 17
(2008). 18

Intrinsic motivations and curiosity

« 45 » **Stojanov** observes that the intrin- 21
sic motivations exhibited by CALM “can 22
be related to low-level physiological drives 23
(hunger, pain-avoidance) with no possibil- 24
ity for development of more sophisticated 25
forms of motivations such as curiosity” 26
(§4). I agree that the motivation system of 27
my model is still far too utilitarian, even 28
though some effort has been made to build 29
an intrinsically motivated agent, which is 30
consistent from the perspective of an em- 31
bodied AI. It is evident that motivation is 32
also linked to the subject's activity itself. 33
Drinking water because we are thirsty is a 34
kind of behavior that can be easily anchored 35
in a biologically-driven explanation. Other 36
behaviors, such as playing checkers, listen- 37 323
ing to music, or writing scientific papers, 38
can hardly be explained by simply referring 39
to physiological needs. 40

« 46 » Nevertheless, CALM imple- 41
ments the notion of curiosity for explora- 42
tory behavior. In naive AI approaches, 43
curiosity usually means doing random ac- 44
tions from time to time. In CALM, there is 45
a measure of exploratory utility that allows 46
the agent to plan actions that may lead to 47
new discoveries, or new knowledge that 48
would enhance its world model. The mech- 49
anism follows two behavioral policies: one 50
to optimize the affective signals, and an- 51
other to optimize the gain of knowledge 52
related to relevant variables. The choice of 53
what action to do depends on the weight- 54
ing of these two policies. 55

column C

column A

Language and reflectivity

« 47 » Thórisson claims that artificial general intelligence cannot be done “without some form of self-programming on the part of the machine, which in turn cannot be achieved without transparency of its operational semantics” (§4) and finishes by questioning the power of the CALM schema formalism “to support models of self... and their ability to support self-inspection” (§6). Similarly, Scott says that “CALM ignores or takes for granted that which is... peculiarly human in human cognition: the ability to communicate and compute using... ‘significant symbols’” (§3).

« 48 » It is evident that those capacities are the notable characteristics of high-level intelligence. But to stay in accordance with constructivist principles, I believe that those abilities emerge as a result of the process of learning and interpreting the experiences on abstract levels. Piaget (1953) suggested that the basic principles regulating intellectual functioning remain unchanged over a

column B

lifetime, and that increasingly refined skills and knowledge result from the gradual complexification of the underlying constructed knowledge structures.

« 49 » To summarize my point, I do not believe that the absence of language or self-investigation represents a particular lack in the mechanism. In fact, the simplicity of the problems faced by CALM for now, as well as its non-existent capacity for creating different layers of abstraction to interpret its experiences, do not allow the agent to have the faculty for doing language or self-awareness.

Conclusion

« 50 » The commentaries pointed out many aspects of my model that can be improved, such as the lack of sensorimotor grounding of the symbolic elements manipulated by CALM and the impossibility to create more abstract levels of knowledge to represent the agent’s experience. I am confident that improved versions of my model will be able to deal with these problems, in

column C

particular extending it to be modular. In this way, the agent will be equipped with more refined sensorimotor apparatus, capable of realizing some pre-processing of signals, coupled with other modules capable of doing some preliminary computing in order to solve some basic sensorimotor problems at a low level, filtering the data that must be sent to the first symbolic modules. Finally, a more sophisticated form of abstraction needs to be incorporated into the algorithm in order to allow the construction of an organized structure of anticipatory modules, acting at different levels of abstraction.

Acknowledgements

I would like to thank all of the commentators for this very enriching discussion, and also thank the editors for this valuable opportunity.

RECEIVED: 5 NOVEMBER 2013

ACCEPTED: 8 NOVEMBER 2013

Combined References

Allen J. W. P. & Bickhard M. H. (2011) Normativity: A crucial kind of emergence. *Human Development* 54: 106–112.

Allen J. W. P. & Bickhard M. H. (2013a) Stepping off the pendulum: Why only an action-based approach can transcend the nativist-empiricist debate. *Cognitive Development* 28: 96–133.

Allen J. W. P. & Bickhard M. H. (2013b) The pendulum still swings. *Cognitive Development* 28: 164–174.

Allen J. W. P. & Bickhard M. H. (2013c) Beyond principles and programs: An action framework for modeling development. *Human Development* 56: 171–177.

Anderson M. (2003) Embodied cognition: A field guide. *Artificial Intelligence* 149(1): 91–130.

Ashby W. R. (1952) *Design for a brain*. First Edition. Chapman & Hall, London.

Barandiaran X. & Moreno A. (2006) On what makes certain dynamical systems cognitive. *Adaptive Behavior* 14(2): 171–185.

column A

Barandiaran X. & Moreno A. (2008) Adaptivity: From metabolism to behavior. *Adaptive Behavior* 16(5): 325–344.

Barsalou L. W. (1999) Perceptual symbol systems. *Behavioral and Brain Sciences* 22: 577–600.

Barsalou L. W. (2008) Grounded cognition. *Annual Review of Psychology* 59: 617–645.

Beer R. D. (1995) A dynamical systems perspective on agent-environment interactions. *Artificial Intelligence* 72: 173–215.

Beer R. D. (2004) Autopoiesis and cognition in the game of life. *Artificial Life* 10(3): 309–326.

Beetz M., Jain D., Mösenlechner L. & Tenorth M. (2010) Towards performing everyday manipulation activities. *Robotics and Autonomous System* 58(9): 1085–1095.

Bellman R. A. (1957) Markovian decision process. *Journal of Mathematics and Mechanics* 6: 679–684.

Bickhard M. H. (2009a) Interactivism: A manifesto. *New Ideas in Psychology* 27: 85–95.

Bickhard M. H. (2009b) The interactivist model. *Synthese* 166(3), 547–591.

column B

Bickhard M. H. & Campbell R. L. (1996) Topologies of learning and development. *New Ideas in Psychology* 14(2): 111–156.

Bickhard M. H. & Richie D. M. (1983) On the nature of representation: A case study of James Gibson’s theory of perception. Praeger Publishers, New York.

Blum A. L. & Langley P. (1997) Selection of relevant features and examples in machine learning. *Artificial Intelligence* 97: 245–271.

Boutilier C., Dearden R. & Goldszmidt M. (2000) Stochastic dynamic programming with factored representations. *Artificial Intelligence* 121(1–2): 49–107.

Bril B., Roux V., Dietrich G. (2005) Stone knapping: Khambhat (India), a unique opportunity? In: Roux V. & Bril B. (eds.) *Stone Knapping: The necessary conditions for a uniquely hominid behaviour*. McDonald Press, Cambridge MA: 53–72.

Brooks R. A. (1991) Intelligence without representation. *Artificial Intelligence* 47: 139–159.

Buisson J.-C. (2004) A rhythm recognition computer program to advocate interactivist perception. *Cognitive Science* 28(1): 75–87.

column C

column A

column B

column C

- 1 Butz M. V. (2002) Anticipatory learning classifier
2 systems. Kluwer, Boston MA.
- 3 Butz M. V. (2008) How and why the brain
4 lays the foundations for a conscious self.
5 Constructivist Foundations 4(1): 1–14 &
6 32–37. Available at [http://www.univie.ac.at/
7 constructivism/journal/4/1/001.butz](http://www.univie.ac.at/constructivism/journal/4/1/001.butz)
- 8 Butz M. V. (2013) Separating goals from behav-
9 ioral control: Implications from learning
10 predictive modularizations. *New Ideas in
11 Psychology* 31: 302–312.
- 12 Byrne R. W. (2005) The maker not the tool: The
13 cognitive significance of great ape manual
14 skills. In: Roux V. & Bril B. (eds.) *Stone
15 Knapping: The necessary conditions for a
16 uniquely hominid behaviour*. McDonald
17 Press, Cambridge MA: 159–169.
- 18 Campbell D. T. (1974) Evolutionary epistemol-
19 ogy. In: Schilpp P. A. (ed.) *The philosophy
20 of Karl Popper*. Open Court, LaSalle IL:
21 413–463.
- 22 Campbell R. L. & Bickhard M. H. (1993) Know-
23 ing levels and the child's understanding of
24 mind. *Behavioral and Brain Sciences* 16(1):
25 33–34.
- 26 Cassandra A. R., Kaelbling L. P. & Littman M.
27 L. (1998) Planning and acting in partially
28 observable stochastic domains. *Artificial
29 Intelligence* 101: 99–134.
- 30 Chaput H. (2004) The constructivist learning
31 architecture. Unpublished PhD Thesis,
32 University of Texas.
- 33 Conant R. C. & Ashby W. R. (1970) Every good
34 regulator of a system must be a model of
35 that system. *International Journal of Systems
36 Science* 1(2): 89–97.
- 37 Crook P. & Hayes G. (2003) Could active percep-
38 tion aid navigation of partially observable
39 grid worlds? In: *Proceedings of the 14th
40 European Conference on Machine Learning
41 (ECML)*. Lecture Notes in Artificial Intel-
42 ligence 2837. Springer-Verlag, Berlin: 72–83.
- 43 Degris T. & Sigaud O. (2010) Factored markov
44 decision processes. In: Buffet O. & Sigaud O.
45 (eds.). *Markov decision processes in artificial
46 intelligence*. Wiley-ISTE, London: 99–125.
- 47 Degris T., Sigaud O. & Wuillemin P.-H. (2006)
48 Learning the structure of factored markov
49 decision processes in reinforcement learning
50 problems. In: *Proceedings of the 23rd Inter-
51 national Conference on Machine Learning*.
52 ACM Press: 257–264.
- 53 Drescher G. (1991) Made-up minds: A construc-
54 tive approach to artificial intelligence. MIT
55 Press, Cambridge MA.
- Feinberg E. A. & Shwartz A. (2002) *Handbook
of Markov decision processes: Methods and
applications*. Kluwer, Norwell MA.
- Foerster H. von (2003) *Understanding under-
standing: Essays on cybernetics and cogni-
tion*. Springer, New York.
- Foner L. N. & Maes P. (1994) Paying atten-
tion to what's important: Using focus of
attention to improve unsupervised learn-
ing. In: Cliff D., Husbands P., Meyer J. &
Wilson S. (eds.) *From animals to animats
3*. *Proceedings of the Third International
Conference on Simulation of Adaptive
Behavior*. MIT Press, Cambridge MA:
256–265.
- Glaserfeld E. von (1974) Piaget and the radical
constructivist epistemology. In: Smock C.
D. & Glaserfeld E. von (eds.) *Epistemology
and education*. Follow Through Publica-
tions, Athens GA: 1–24. Available at [http://
www.vonglaserfeld.com/034](http://www.vonglaserfeld.com/034)
- Glaserfeld E. von (1979) Cybernetics, experi-
ence, and the concept of self. In: Ozer
M. N. (ed.) *A cybernetic approach to the
assessment of children: Toward a more
humane use of human beings*. Westview
Press: Boulder, CO, pp. 67–113. Available at
[http://www.vonglaserfeld.com/056/
Glaserfeld E. von \(2005\) 30 years of radical
constructivism. Constructivist Foundations
1\(1\): 9–12. Available at \[http://www.univie.
ac.at/constructivism/journal/1/1/009.
glaserfeld\]\(http://www.univie.ac.at/constructivism/journal/1/1/009.glaserfeld\)](http://www.vonglaserfeld.com/056/)
- Goldstein J. (1999) Emergence as a construct:
History and issues. *Emergence: Complexity
and Organization* 1: 49–72.
- Guerin F. (2011) Learning like baby: A survey
of AI approaches. *Knowledge Engineering
Review* 26(2): 209–236.
- Guestrin C., Koller D. & Parr R. (2001) Solving
factored POMDPs with linear value func-
tions. In: *Proceedings of the Workshop on
Planning under Uncertainty and Incom-
plete Information*. Morgan Kaufmann, San
Francisco CA: 67–75.
- Guestrin C., Koller D., Parr R. & Venkatara-
man S. (2003) Efficient solution algorithms
for factored MDPs. *Journal of Artificial
Intelligence Research* 19: 399–468.
- Hansen E. A. & Feng Z. (2000) Dynamic
programming for POMDPs using a factored
state representation. In: *Proceedings of the
Fifth International Conference on Artificial
Intelligence, Planning and Scheduling*.
AAAI Press, Menlo Park CA:130–139.
- Hauskrecht M. (2000) Value-function approxi-
mations for partially observable Markov
decision processes. *Journal of Artificial Intel-
ligence Research* 13: 33–94.
- Hoffmann J. (2003) Anticipatory behavioral
control. In: Butz M. V., Sigaud O. & Gérard
P. (eds.) *Anticipatory behavior in adaptive
learning systems: Foundations, theories, and
systems*. Springer, Berlin: 44–65.
- Holmes M. & Isbell C. (2005) Schema learning:
Experience-based construction of predictive
action models. *Advances in Neural Informa-
tion Processing Systems* 17: 562–585.
- Holmes M. & Isbell C. (2006) Looping suffix
tree-based inference of partially observable
hidden state. In: Cohen W. W. & Moore A.
(eds.) *Proceedings of the 23rd International
Conference on Machine Learning (ICML
2006)*. ACM Press, New York: 409–416.
- Hommel B., Müsseler J., Aschersleben G. &
Prinz W. (2001) The theory of event coding
(TEC): A framework for perception and ac-
tion planning. *Behavioral and Brain Sciences*
24: 849–878.
- Indurkha B. (1992) *Metaphor and cogni-
tion: An interactionist approach*. Kluwer,
Dordrecht.
- Joas H. (1993) American pragmatism and Ger-
man thought: A history of misunderstand-
ings. In: Joas H., *Pragmatism and social
theory*. University of Chicago Press, Chicago:
94–121.
- Jonsson A. & Barto A. (2005) A causal approach
to hierarchical decomposition of factored
MDPs. In: De Raedt L. & Wrobel S.(eds.)
*Proceedings of the 22nd International Con-
ference on Machine Learning (ICML 2005)*.
ACM Press, New York: 401–408.
- Kitchener F. K. (1986) *Piaget's theory of knowl-
edge*. Yale University Press, New Haven.
- Maturana H. R. & Varela F. J. (1980) *Autopoiesis
and cognition*. Reidel, Dordrecht.
- Minsky M. (1988) *The society of mind*. Simon &
Schuster, New York.
- Modayil J., White A. & Sutton R. S. (2012)
Multi-timescale nexting in a reinforcement
learning robot. In: Ziemke T., Balkenius
C. & Hallam J. (eds.) *Proceedings of the
Conference on the Simulation of Adaptive
behavior (SAB'12)*. Springer, Heidelberg:
299–309.
- Morrison C., Oates T. & King G. (2001)
Grounding the unobservable in the observ-
able: The role and representation of hidden
state in concept formation and refinement.

column A

column B

column C

column A

- 1 In: Working notes of AAAI Spring Symposium. AAAI Press, Menlo Park CA: 45–49.
- 2
- 3 Nivel E. & Thórisson K. R. (2013) Towards
- 4 a programming paradigm for control
- 5 systems with high levels of existential
- 6 autonomy. In: Kühnberger K.-U., Rudolph S. & Wang P. (eds.) Artificial General Intelligence. Proceedings of the Sixth International Conference, AGI-13, Beijing, China, July/August 2013. Springer, Berlin:
- 7 78–87.
- 8
- 9 Nivel E., Thórisson K. R., Dindo H., Pezzulo G., Rodriguez M., Corbato C., Steunebrink B., Ognibene D., Chella A., Schmidhuber J., Sanz R. & Helgason H. P. (2013) Autocatalytic endogenous reflective architecture. Reykjavik University School of Computer Science Technical Report, RUTR-SCS13002.
- 10
- 11 Parisi D. (2004) Internal robotics. *Connection Science* 16(4): 325–338.
- 12
- 13 Pask G. (1981) Organisational closure of potentially conscious systems In: Zelany M. (ed.) *Autopoiesis*. North Holland Elsevier: New York: 265–307.
- 14
- 15 Pask G., Scott B. & Kallikourdis D. (1973) A theory of conversations and individuals (exemplified by the learning process on CASTE). *International Journal of Man-Machine Studies* 5: 443–566.
- 16
- 17 Pastra K. & Aloimonos Y. (2012) The minimalist grammar of action. *Philosophical Transactions of the Royal Society B: Biological Sciences* 367: 103–117.
- 18
- 19 Perotto F. S. (2010) Un mécanisme constructiviste d'apprentissage automatique d'anticipations pour des agents artificiels situés. Unpublished Ph.D Thesis at the Institut National Polytechnique de Toulouse, France.
- 20
- 21 Piaget J. (1953) *The origin of intelligence in the child*. Routledge & Kegan Paul. London.
- 22
- 23 Piaget J. (1954) *The construction of reality in the child*. Routledge & Kegan Paul, London. Originally published as: Piaget J. (1937) *La construction du réel chez l'enfant*. Delachaux et Niestlé, Neuchâtel.
- 24
- 25 Poupart P. & Boutilier C. (2004) VDCBPI: An approximate scalable algorithm for large scale POMDPs. In: Proceedings of the 17th Advances in Neural Information Processing Systems (NIPS). MIT Press, Cambridge MA: 1081–1088.
- 26
- 27 Prigogine I. (1980) *From being to becoming*. Freeman, San Francisco CA.

column A

column B

- Puterman M. L. (1994) *Markov decision processes: Discrete stochastic dynamic programming*. Wiley, New York.
- Rivest R. & Schapire R. (1994) Diversity-based inference of finite automata. *ACM Journal* 43(3): 555–589.
- Schachner W. (ed.) (1996) *Constructivism, interactionism and their applications*. CEPIAG's Symposium Proceeding. CEPIAG, Geneva.
- Schachner W., Real del Sarte O. & León C. (1999) Valore, indice e rappresentazione nella guida dei valori elementari. *Metis* 1: 7–22.
- Scott B. (2007) The co-emergence of parts and wholes in psychological individuation. *Constructivist Foundations* 2(2–3): 65–71. Available at <http://www.univie.ac.at/constructivism/journal/2/2-3/065.scott>
- Scott B. & Shurville S. (2011) What is a symbol? *Kybernetes* 48(1/2): 12–22.
- Searle J. (1980) Minds, brains, and programs. *Behavioural and Brain Sciences* 3: 417–424.
- Shani G., Brafman R. I. & Shimony S. E. (2005) Model-based online learning of POMDPs. In: Proceedings of the 16th European Conference on Machine Learning (ECML). Springer, Berlin: 353–364.
- Shani G., Poupart P., Brafman R. I. & Shimony S. E. (2008) Efficient ADD operations for point-based algorithms. In: Proceedings of the 8th International Conference on Automated Planning and Scheduling (ICAPS). AAAI Press, Menlo Park CA: 330–337.
- Sigaud O., Butz M. V., Kozlova O. & Meyer C. (2009) Anticipatory learning classifier systems and factored reinforcement learning. In: Pezzulo G., Butz M. V., Sigaud O. & Baldassarre G. (eds.) *Anticipatory behavior in adaptive learning systems: From psychological theories to artificial cognitive systems*. Springer, Berlin: 321–333.
- Sim H. S., Kim K.-E., Kim J. H., Chang D.-S. & Koo M.-W. (2008) Symbolic heuristic search value iteration for factored POMDPs. In: Proceedings of the 23rd National Conference on Artificial Intelligence. AAAI Press, Menlo Park CA: 1088–1093.
- Singh S., James M. & Rudary M. (2004) Predictive state representations: A new theory for modeling dynamical systems. In Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence. AUAI Press, Banff: 512–519.
- Singh S., Littman M., Jong N., Pardoe D & Stone P. (2003) Learning predictive state

column B

column C

- representations. In: Proceedings of the 20th International Conference on Machine Learning (ICML). AAAI Press, Menlo Park CA: 712–719.
- Stojanov G. (2009) History of usage of Piaget's theory of cognitive development in AI and robotics: A look backwards for a step forwards. In: Cañamero L., Oudeyer P.-Y. & Balkenius C. (eds.) Proceedings of the Ninth International Conference on Epigenetic Robotics. LUCS, Lund: 243–245. Available at <http://www.lucs.lu.se/LUCS/146/>
- Stojanov G., Bozinovski S. & Trajkovski G. (1997) Interactionist-expectative view on agency and learning. *IMACS Journal for Mathematics and Computers in Simulation* 44: 295–310.
- Stojanov G. & Kulakov A. (2011) Modeling attention within a complete cognitive architecture. In: Roda C. (ed.) *Attention support in digital environments*. Cambridge University Press, Cambridge: 210–241.
- Stojanov G., Kulakov A. & Clauzel D. (2006) On curiosity in intelligent robotic systems. In: Proceedings AAAI 2006 Fall Symposium on Interaction and Emergent Phenomena in Societies of Agents, Arlington, Virginia. AAAI Press, Cambridge: 183–189.
- Stolzmann W. (2000) An Introduction to anticipatory classifier systems. In: Lanzi P. L., Stolzmann W. & Wilson S. W. (eds.) *Learning classifier systems: From foundations to applications*. Springer, Berlin: 175–194.
- Stoytchev A. (2009) Some basic principles of developmental robotics. *IEEE Transactions on Autonomous Mental Development* 1(2): 122–130.
- Strehl A. L., Diuk C. & Littman M. L. (2007) Efficient structure learning in factored-state MDPs. In: Proceedings of the 22nd National Conference on Artificial Intelligence, AAAI Press, Menlo Park CA: 645–650.
- Suchman L. A. (1987) *Plans and situated actions: The problem of human-machine communication*. Cambridge University Press, Cambridge.
- Sutton R. S. & Barto A. G. (1998) *Reinforcement learning: An introduction*. MIT Press, Cambridge MA.
- Sutton R. S., Modayil J., Delp M., Degris T., Pilarski P. M., White A. & Precup D. (2011) Horde: A scalable real-time architecture for learning knowledge from unsupervised sensorimotor interaction. In: Proceedings of the Tenth International Conference on

column C

column A	column B	column C
1 Autonomous Agents and Multiagent Systems	and its pernicious effect on artificial intelli-	Wörgötter F., Aksoy E. E., Krüger N., Piater J.,
2 (AAMAS'11), Volume 2. IFAAMAS, Taipei:	gence research. In: Abdel-Fattah A. H. M. &	Ude A. & Tamosiunaite M. (2013) A simple
3 761–776.	Kühnberger K.-U. (eds.) Proceedings of the	ontology of manipulation actions based on
4 Sutton R. S., Precup D. & Singh S. (1999)	Workshop Formalizing Mechanisms for Ar-	hand-object relations. IEEE Transactions
5 Between MDPs and semi-MDPs: A frame-	tificial General Intelligence and Cognition	on Autonomous Mental Development 5:
6 work for temporal abstraction in reinforce-	(Formal MAGiC). Institute of Cognitive	117–134.
7 ment learning. Artificial Intelligence 112:	Science, Osnabrück: 31–35.	
8 181–211.	Thornton C. (2003) Indirect sensing through	Yavuz A. & Davenport D. (1997) PAL: A con-
9 Tani J. (1996) Model-based learning for mobile	abstractive learning. Intelligent Data Analy-	structivist model of cognitive activity. In:
10 robot navigation from the dynamical	sis 7(3): 1–16.	Riegler A., Peschl M. F. & von Stein A. (eds.)
11 systems. Perspective. IEEE Transactions	Tortora G. J. & Derrickson B. H. (2012)	Proceedings of the international conference
12 on Systems, Man, and Cybernetics, Part	Principles of anatomy and physiology. 13th	“New Trends in Cognitive Science: Does
13 B: Cybernetics, Special Issue on learning	Edition. Wiley, New York.	Representation Need Reality?” Vienna,
14 autonomous robots 26(3): 421–436.	Varela F., Thompson E. & Rosch E. (1991) The	Austria.
15 Thórisson K. R. (2012) A new constructivist	embodied mind: Cognitive science and	Ziemke T. (2003) What’s that thing called em-
16 AI: From manual construction to self-con-	human experience. MIT Press, Cambridge	bodiment? In: Alterman R. & Kirsh D. (eds.)
17 structutive systems. In: Wang P. & Goertzel B.	MA.	Proceedings of the 25th Annual Meeting
18 (eds.) Theoretical foundations of artificial	Wilson R. & Clark A. (2008) How to situate	of the Cognitive Science Society. Lawrence
19 general intelligence. Atlantis Press, Paris:	cognition: Letting nature take its course. In:	Erlbaum, Mahwah NJ: 1305–1310.
20 145–171.	Aydene M. & Robbins P. (eds.). Cambridge	
21 Thórisson K. R. (2013) Reductio ad absurdum:	handbook of situated cognition. Cambridge	
22 On oversimplification in computer science	University Press, New York.	
23		
24		
25		
26		
27		
28		
29		
30		
31		
32		
33		
34		
35		
36		
37		
38		
39		
40		
41		
42		
43		
44		
45		
46		
47		
48		
49		
50		
51		
52		
53		
54		
55		