Innate Theories as a Basis for Autonomous Mental Development

Thomas C. Henderson
University of Utah

Xiuyi Fan
Australian Centre for Field Robotics

Aniesha Alford
North Carolina A& T

Edward Grant
North Carolina State University
Elaine Cohen
University of Utah

UUCS-09-004

School of Computing
University of Utah
Salt Lake City, UT 84112 USA

9 October 2009

Abstract

Sloman (in robotics), Chomsky and Pinker (in natural language), and others, e.g., Rosenberg (in human cooperative behavior) have proposed that some abstract theories relevant to cognitive activity are encoded genetically in humans. The biological advantages of this are (1) to reduce the learning time for acquisition of specific contextual models (e.g., from a language community; appropriate physics, etc.), and (2) to allow the determination of true statements about the world beyond those immediately available from direct experience. We

believe that this hypothesis is a strong paradigm for the autonomous mental development of artificial cognitive agents and we give specific examples and propose a theoretical and experimental framework for this. In particular, we show that knowledge and exploitation of symmetry can lead to greatly reduced reinforcement learning times on a selected set of problems.

1 Introduction

Cognitive systems perceive, deliberate and act in unstructured environments, and the development of effective mental abilities is a longstanding goal of the AI and intelligent systems communities. The major approaches are the *cognitivist* (physical symbol systems) and *emergent* (dynamical systems) paradigms. For a detailed review of the relevant characteristics of cognitive systems and how these two approaches differ, see [9]. Basically, cognitivists maintain that patterns of symbol tokens are manipulated syntactically, and through percept-symbol associations, perception is achieved as abstract symbol representations and actions are causal consequences of symbol manipulation. In contrast, emergent systems are concurrent, self-organizing networks with a global system state representation which is semantically grounded through skill construction, where perception is a response to system perturbation and action is a perturbation of the environment by the system. The emergent approach searches the space of closed-loop controllers to build higher-level behavior sequences out of lower ones so as to allow a broader set of affordances in terms of the sensorimotor data stream. We propose to combine these approaches in order to exploit abstraction and specific domain theories to overcome that complexity. The hypothesis is:

The Domain Theory Hypothesis: Semantic cognitive content may be effectively discovered by restricting controller solutions to be models of specific domain theories intrinsic to the cognitive architecture.

Sloman [6, 8, 7] has argued for this from a philosophical point of view, while Chomsky [2] and Pinker [4] have explored universal structures for human natural language, and Rosenberg [5] explores the genetic evidence for cooperative behavior among humans. We study the hypothesis in the context of some standard AI and robotics problems. In particular, we consider here the role that a theory of symmetry can play in various learning scenarios. When symmetry can be exploited in reinforcement learning, the time to learn the solution to the task should be proportional to the size of the set of asymmetric states (note that this may be characterized in terms of the quotient space of the associated group where it exists). Figure 1 shows the cognitive architecture for this approach.

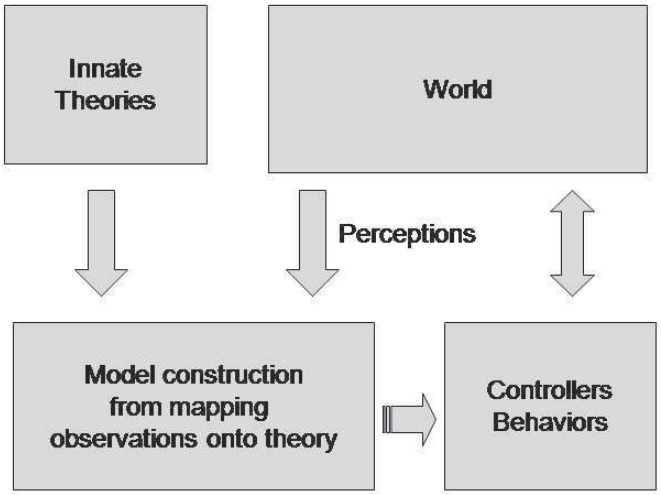

Figure 1: Innate Theory based Cognitive Architecture.

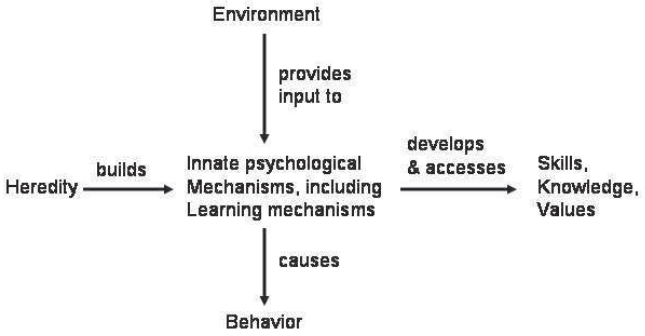

Figure 2: Pinker's Computational Modules Cognitive Architecture [4].

Pinker (p. 423) previously proposed the schema shown in Figure 2 for innate computational modules in humans; he also outlined the following tests for possible computational modules (what we call theories) in humans and gives some examples (pp. 436-438):

1. Does the theory help solve a problem that our ancestors faced in their environment (biological anthropology)?
2. When children solve problems for which mental modules exist, they should know things they have not been taught.
3. Neuroscience should discover that the brain tissue computing the problem has some kind of physiological cohesiveness (tissue or subsystem).

Pinker also lists some possible modules:

- Intuitive mechanisms: knowledge of the motions, forces and deformations that objects undergo.
- Intuitive biology: understanding how plants and animals work.
- Number
- Mental maps for large territories
- Habitat selection: seeking of safe, information-rich, productive environments, generally savannah-like.
- Danger, including the emotions of fear and caution, phobias for stimuli such as heights, confinement, risk, social encounters, venomous and predatory animals, and a motive to learn the circumstances in which each is harmless.
- Food: what is good to eat.
- Contamination, including the emotion of disgust, reactions to certain things that seem inherently disgusting, and intuition about contagion and disease.
- Monitoring of current well-being, including emotions of happiness and sadness, and moods of contentment and restlessness.
- Intuitive psychology: predicting other people's behavior from their beliefs and desires.
- A mental Rolodex: a database of individuals, with blanks for kinship, status or rank, history of exchange of favors, and inherent skills and strengths, plus criteria that valuate which trait.
- Self-concept: gathering and organizing information about one's value to other people, and packaging it for others.
- Justice: sense of rights, obligations, and deserts, including the emotions of anger and revenge.
- Kinship: including nepotism and allocations of parental effort.
- Mating: including feelings of sexual attraction, love and intentions of fidelity and desertion.

In this paper, we explore theories of knowledge of the motions, forces and deformations that objects undergo.

Moreover, we believe that a hierarchical structure for these modules is more suitable, that is, some are more basic than others. Another related question is whether there is a specific set of identifiable domains to which theories apply. Of course, many questions arise. Are these domains predefined by the environment or derived during developmental stages of the organism? Are innate theories the same among all humans? Do innate theories determine the difference between human and non-human? Are theories ever modifiable? (Do they evolve?) Do theories come in various types? i.e., are some more specific and others more general? How are theories learned? How are theories, if they exist, represented? Symbolically? Analogically? Innate knowledge should be reflected in the structure of the brain and body. For example, knowledge of symmetry may first be realized by observing facts that humans have bilateral symmetry: two hands, two legs, two eyes, etc.

The *Domain Theory* predicates:

1. A representation of an innate theory and inference rules for the theory.
2. A perceptual mechanism to determine elements of a set and operators on the set.
3. A mechanism to determine that the set and its operators are a model of the innate theory.
4. Mechanisms to allow the exploitation of the model in learning and belief construction.

As an example, consider group theory. Given a set, S, and operator, $+$, there are four axioms:

1. **Closure**: $a, b \in S \Rightarrow a + b \in S$
2. **Associativity**: $a + (b + c) = (a + b) + c$
3. **Identity element**: $\exists e \in S \ni a \in S \Rightarrow a + e = a$
4. **Inverse**: $\forall a \in S \exists a^{-1} \in S \ni a + a^{-1} = e$

For a detailed discussion of a logical axiomatization of group theory, see [3]. Hamilton also states that the first order theory of Abelian groups is recursively decidable.

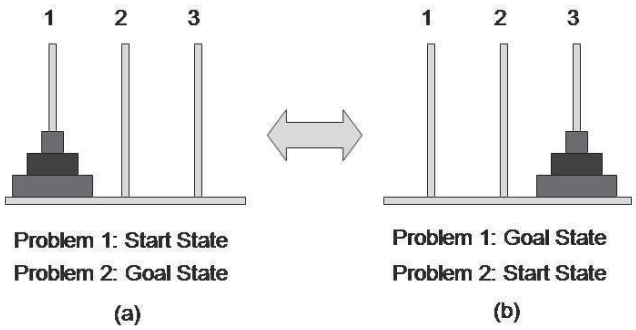

Figure 3: Towers of Hanoi Problem. Moving one ring at a time from its location (1, 2, or 3) to one of 1, 2 or 3 such that a larger ring is never placed on top of a smaller ring, transform the state shown in (a) to that shown in (b) or vice versa).

The set of rotations of rigid objects in the plane consisting of the two rotations by 0 and π radians (call these E and R) and the operation of applying two rotations in sequence constitute a group. This can be seen since (1) $E^2 = E$, $ER = RE = E$, $R^2 = E$, (2) a sequence of these operations is associative, (3) the identity is E, and (4) $E^{-1} = E$ and $R^{-1} = R$. Thus, the rotations constitute a model of group theory. In this paper we do not address how this model is discovered, but identify this as an important research topic.

In the remaining sections of the paper we discuss a couple of examples of exploiting a model of symmetry to make a learning task more efficient. In this case, we suppose that specific states are known to be symmetric.

2 Simple Reinforcement Learning Example

Consider the Towers of Hanoi problem (see Figure 3) where we consider learning the two solutions: from left tower to right and vice versa. For this example, we exploit the dihedral symmetry group D_1. Dihedral groups $D_1, D_2, \ldots$ where D_n (of order $2n$) consist of the rotations in C_n (the cyclic groups $C_1, C_2, \ldots C_n$ consist of all rotations about a fixed point by multiples of the angle 360 degrees) together with reflections in n axes that pass through the fixed point. C_1 is the trivial group containing only the identity operation, which occurs when the figure has no symmetry at all. D_1 is the 2-element group containing the identity operation and a single reflection, which occurs when the figure has only a single axis of bilateral symmetry, as is the case for the poles in the Towers of Hanoi problem. For more details see [1] and Wikipedia.

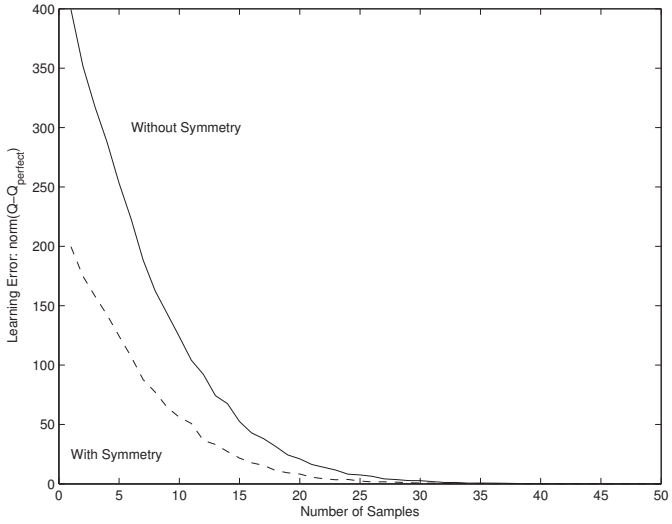

Figure 4: Towers of Hanoi Learning Rates with and without Symmetry.

We can take advantage of symmetry during reinforcement learning as follows:

1. Determine the set of symmetric states.
2. During selection of a starting state in the RL algorithm, always choose the same representative of the respective symmetric states.
3. When assigning credit during the Q table update step, update symmetric states, too.

When symmetry is treated this way, the time to learning is reduced as shown in Figure 4. In this figure the abscissa gives the number of samples used to learn (from 0 to 50), where each sample result is averaged over 100 trials, while the ordinate is the measure of learning error (i.e., the distance between the perfect solution and the learned solution). Learning results without symmetry are shown by the solid (blue) line, whereas learning results using symmetry are shown with the dashed (red) line. The symmetry is exploited during reinforcement learning by updating the Q table entry for $(sym_state, sym_next_state)$ whenever $(state, next_state)$ is updated. Without symmetry, reinforcement learning must be applied to two distinct reward tables (one for going left to right and the other for going right to left). As can be seen, the learning time is reduced about 50% when symmetry is exploited and that this corresponds to the number of symmetric states.

In order to determine that symmetry applies, the cognitive agent must be able to map observed essential aspects of the problem domain onto the appropriate theory which in this

case is group theory for D_1. The required mapping must take the left and right poles into each other (reflection about the y axis). The rings are a secondary consideration, but the situation can be analyzed as follows. Let (P_1, P_2, P_3) represent the pole location of the large, medium and small rings in that order, where the left pole is 1, the middle pole is 2 and the right pole is 3. Then the states, S_i are given by:

```
S1: (1,1,1)        S15: (3,2,3)
S2: (1,1,2)        S16: (2,2,1)
S3: (1,1,3)        S17: (2,1,2)
S4: (1,3,2)        S18: (3,3,1)
S5: (1,2,3)        S19: (3,1,3)
S6: (1,3,1)        S20: (2,2,2)
S7: (1,3,3)        S21: (2,2,3)
S8: (1,2,1)        S22: (2,1,3)
S9: (1,2,2)        S23:(2,1,1)
S10: (2,3,3)       S24: (3,3,3)
S11: (3,2,2)       S25: (3,3,2)
S12: (2,3,1)       S26: (3,1,1)
S13: (2,2,2)       S27: (3,1,2)
S14: (3,2,1)
```

The state transitions are given by:

State	Next States	State	Next States
1	2,3	15	11,14,19
2	1,3,4	16	12,20,21
3	1,2,5	17	13,22,23
4	2,6,7	18	14,24,25
5	3,8,9	19	15,26,27
6	4,7,8	20	16,21
7	4,6,10	21	16,20,22
8	5,6,9	22	17,21,23
9	5,8,11	23	17,22,26
10	7,12,13	24	18,25
11	9,14,15	25	18,24,27
12	10,13,16	26	19,23,27
13	10,12,17	27	19,25,26
14	11,15,18		

Finally, the symmetric states are (where $\perp$ represents *is symmetric*):

$$S = \{S_1 \perp S_{24}, S_2 \perp S_{25}, S_3 \perp S_{18}, S_4 \perp S_{27}, S_5 \perp S_{14}, S_6 \perp S_{19}, S_7 \perp S_{26}, S_8 \perp S_{15}, S_9 \perp S_{11}, S_{10} \perp S_{23}, S_{12} \perp S_{22}, S_{13} \perp S_{17}, S_{16} \perp S_{21}, S_{20} \perp S_{20}\}$$

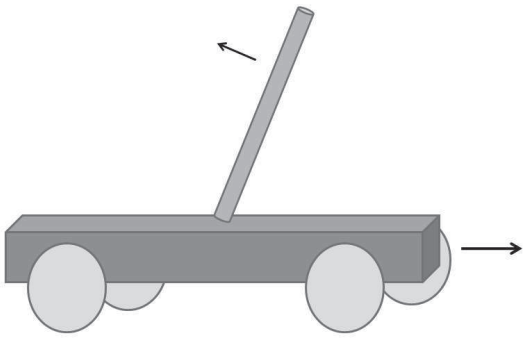

Figure 5: Cart-Pole Problem.

3 Pole-Cart Problem

The previous example involved discrete states and operators, and we now consider a continuous state problem. The goal is to balance a 1 DOF pole (a rotary joint connects the pole to a mobile cart) atop a cart which can move in $\pm x$ (see Figure 5). Motion is achieved by applying a bang-bang control to the cart by means of a force of magnitude F applied in either the plus or minus x direction. A state is defined as a specific set of values for $(x, \dot{x}, \theta, \dot{\theta})$, where x is the position of the cart (and the base of the pole), $\dot{x}$ is the velocity of the cart, θ is the angle of the pole from vertical (clockwise being positive rotation), and $\dot{\theta}$ the angular velocity of the pole.

The goal is to use reinforcement learning to find the control law which maintains θ as near 0 as possible. Symmetric states are given by:

$$(x, \dot{x}, \theta, \dot{\theta}) \perp -(x, \dot{x}, \theta, \dot{\theta})$$

From any given state, the possible actions are to apply a force of $-F$ or F. Thus, just as with the Towers of Hanoi problem, symmetry can be used to update states and learn twice as fast as without symmetry.

When symmetry is treated this way, the time to learning is reduced as shown in Figure 6.

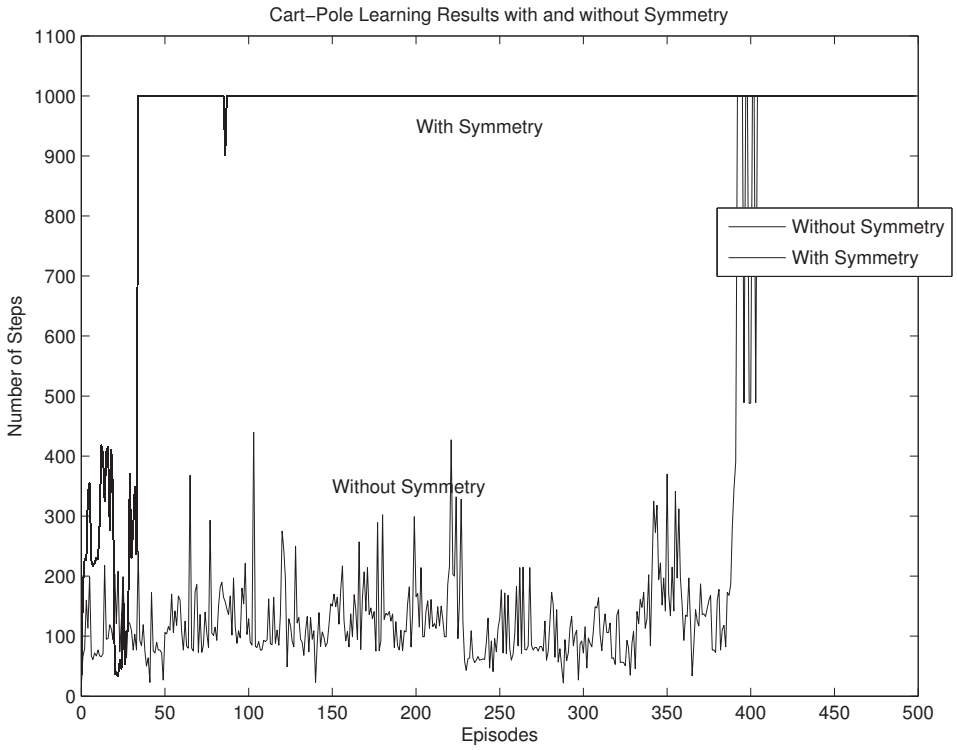

Figure 6: Cart-Pole Learning Rates with and without Symmetry. Note that with symmetry, the pole is kept upright for 1,000 steps after about 30 episodes.

4 Discussion

We have shown here that the exploitation of simple theory can improve the learning rate of a cognitive agent. However, even broader advantages arise. For example, given an appropriate theory, true statements can be discovered (by syntactic inference or semantic truth analysis) which are beyond the phenomena observed by the agent. Moreover, such theorems will be considered true and need not have a probabilistic structure. Where competing theories exist in a single cognitive agent (perhaps due to alternative axioms), it is possible to represent and entertain contradictory conclusions (e.g., believe A and $\neg A$ simultaneously), but without falling into inconsistency since the source of the difference can be referenced.

The use of a simple Boolean algebra or sentential calculus could piggy-back off of natural language structures in humans so that these are identified as the set of sentences (propositions) to which the Boolean operators apply. Then for certain models, the agent would have complete and consistent theories available for the determination of the truthfulness of statements.

Another key question that arises is for which domains such theories might exist, as posed by Pinker above. We believe that this gives rise to a vigorous research agenda:

- What theories are appropriate for which domains?
- How are theories represented in the cognitive agent?
- How are observations mapped to the appropriate theory (i.e., how are models created)?
- How are such models exploited to improve learning?
- How are such models exploited to arrive at new knowledge?

These questions can be studied in animals as well as artificial cognitive agents, and give rise to deep questions about brain form and function, as well as to the possible genetic coding of innate theories.

These also give rise to certain predictions concerning the exploitation of symmetry in humans and some requirements on artificial agents:

- Agents should be able to perceive symmetry. That is, visual analysis mechanisms should exist which respond to visually symmetric stimuli.

- Mechanisms should exist to exploit symmetric relations during learning in certain organisms and not others.
- Using symmetry should reduce the size of the search space during learning (e.g., by about half for the Towers of Hanoi problem).

We are currently exploring the use of innate theories in broader cognitive control systems: autonomous mobile robots and vehicles, as well as in cognitive sensor actuator networks.

References

[1] B. Baumslag and B. Chandler. *Group Theory*. McGraw-Hill Book Co, New York, NY, 1968.

[2] N. Chomsky. Biolinguistic Explorations: Design, Development, Evolution. *International Journal of Philosophical Studies*, 15(1):1–21, March 2007.

[3] A.G. Hamilton. *Logic for Mathematicians*. Cambridge University Press, Cambridge, UK, 1988.

[4] Steven Pinker. *The Language Instinct*. Harper Collins, NY, NY, 1994.

[5] A. Rosenberg. Will Genomics Do More for Metaphysics than Locke? In P. Achinstein, editor, *Scientific Evidence*, pages 186–205, Baltimore, MD, 2005. The Johns Hopkins University Press.

[6] A. Sloman. Architectural and Representational Requirements for Seeing Processes and Affordances. In D. Heinke and E. Mavritsaki, editors, *Computational Modelling in Behavioural Neuroscience: Closing the Gap between NeuroPhysiology and Behaviour*, London, UK, 2008. Psychology Press.

[7] A. Sloman. Some Requirements for Human-like Robots: Why the Recent Over-emphasis on Embodiment has Held up Progress. In B. Sendhoff, E. Koerner, O. Sporns, H. Ritter, and K. Doya, editors, *Creating Brain-like Intelligence*, Berlin, Germany, 2009. Springer-Verlag.

[8] Aaron Sloman. The Well Designed Young Mathematician. *Artificial Intelligence*, 172(18):2015–2034, December 2008.

[9] D. Vernon, G. Metta, and G. Sandini. A Survey of Artificial Cognitive Systems: Implications for the Autonomous Development of Mental Capabilities in Computational

Agents. *IEEE Transactions on Evolutionary Computing, Special Issue on Autonomous Mental Development*, to appear, 2008.