The Role of Symmetry in Structural Bootstrapping

Thomas C. Henderson, Hongchang Peng, Nikhil Deshpande, and Edward Grant

Abstract Robots will play an increasing role in society as they are deployed as co-workers, co-protectors and co-inhabitants among humans, and it is crucial that their knowledge be acquired efficiently and be as correct and effective as possible, and permit planning and rational behavior selection. To achieve this, robot knowledge needs to span several levels (from perception-action processes to concepts). There are several major issues to be addressed in order to achieve this. First, the fundamentally different paradigms for cognitive robotics include Turing machines, neural networks and dynamical systems. Each has starkly different views on what constitutes a concept or perception-actuation mechanism. There may also be a developmental stage in which the robot agent discovers, through self-exploration, its own sensorimotor structure; its representations are then intrinsic to its embodiment. All these factors make well-founded conceptualization difficult. Given the scale of this problem, human specification of cognitive content seems precluded, but general learning structures and dynamical systems approaches may produce idiosyncratic results. What is clear is that constraints from these issues must inform any knowledge representation methodology. We describe a set of core symmetry-based representations and processes for structural bootstrapping which will permit this deeper kind of knowledge representation, and specifically address how low-level symmetry detectors in 1-D, 2-D, and 3-D data can help solve the sensorimotor reconstruction problem.

Thomas C. Henderson
School of Computing, University of Utah, e-mail: tch@cs.utah.edu

Hongchang Peng
School of Computing, University of Utah e-mail: naytech.nstc@gmail.com

Nikhil Deshpande
Electrical and Computer Engineering, North Carolina State University e-mail: nadeshpa@ncsu.edu

Edward Grant
Electrical and Computer Engineering, North Carolina State University e-mail: egrant@ncsu.edu

1 Introduction

Physical robot systems have been steadily improving for many years now in terms of their capabilities, robustness, compliance, etc., and there is a strong push to introduce these systems into human environments as cooperative agents to assist people in their daily activities. A major roadblock to this goal is the lack of strong and robust cognitive abilities in robots, and more specifically inadequate knowledge acquisition, representation and manipulation. Robots need various kinds of knowledge to perform effectively in real applications, and the current approaches to providing that knowledge are to: (1) have the robot learn from scratch, (2) spoon feed the knowledge by human programming, or (3) have the robot use the web to find the appropriate knowledge.

Our goal is to explore the use of symmetry analysis as a basis for the semantic grounding of sensorimotor affordance knowledge; this includes symmetry detection in signals, symmetry parsing in knowledge representation, and symmetry exploitation in structural bootstrapping and knowledge sharing. We are working in close cooperation with our colleagues (Profs. R. Dillmann and T. Asfour at the Karlsruhe Institute of Technology) involved in the European Union Xperience project (http://www.xperience.org/). The overview they give of the Xperience project is:

> Current research in enactive, embodied cognition is built on two central ideas: 1) Physical interaction with and exploration of the world allows an agent to acquire and extend intrinsically grounded, cognitive representations and, 2) representations built from such interactions are much better adapted to guiding behavior than human crafted rules or control logic. The Xperience project will address this problem by *structural bootstrapping*, an idea taken from child language acquisition research. Structural bootstrapping is a method of building generative models, leveraging existing experience to predict unexplored action effects and to focus the hypothesis space for learning novel concepts. This developmental approach enables rapid generalization and acquisition of new knowledge and skills from little additional training data.

This gives us a larger context within which to test our hypothesis that symmetry-based structuring of knowledge provides a more robust semantic basis than current methods, and in particular symmetry as applied to: the acquisition of affordances from signals, representation as Object-Action Complexes (OACs provide a framework for modeling actions and their effects), and exploitation in generative action discovery (structural bootstrapping). We aim to directly compare our results with those of the Xperience team in terms of specific cognitive benchmarks.

In order to achieve structural bootstrapping, it is necessary to have a structure with identifiable elements and relations which can be exploited to attain greater knowledge. This can occur by broadening the domain of application of knowledge (e.g., from knowing how to open a specific bottle to knowing how to open any bottle of that type), or by recognizing an instance of a more general structure (e.g., recognizing a rotary joint). We describe here how both can be achieved in the context of the sensorimotor reconstruction problem.

In order to solve this problem, one form of knowledge of particular interest is self-knowledge about the robot's own structure and capabilities; this includes sensors, actuators, kinematic and dynamic structure, energy consumption and replenishment, and computational capabilities (speed, space, parallel processing, signal processing, internet connectivity, etc.). This provides a basis for knowledge of affordances in the external world, i.e., the recognition of entities appropriate for the performance of a task. Working knowledge is also needed for the interactions between the robot and the environment for both physical actions and social interactions (communication, cooperation, empathy, etc.). Of course, a robot will also need to be able to understand and formulate goals and the plans necessary to achieve those goals, but we do not address this aspect of cognition here.

Although the specific knowledge required by a robot will depend on the particular application domain (e.g., security, surgery, manufacturing, home services, etc.), there is a need for fundamental mechanisms which allow each individual robot to obtain the requisite knowledge. Our view is that current methods are too brittle and do not scale very well, and that a new approach to knowledge acquisition and sharing is necessary. This new approach should provide firm semantic grounding in the real world, provide for robust dynamic performance in real-time environments and allow for communication of acquired knowledge in a broad community of other robots and agents, including humans. We thus formulate the following hypothesis:

> Robot affordance knowledge acquisition and sharing can be enabled by means of a common sensorimotor semantics which is provided by a set of group symmetry theories embedded a priori in each robot. These theories inform the production of structural representations of sensorimotor processes which, in turn, permit structural bootstrapping.

2 Robot Knowledge Sharing

Much previous work on robot knowledge sharing has focused on things like multimedia databases ([9, 14, 15]), ontologies ([1, 2, 40, 42, 43, 47]), etc. For example, our work on RobotShare envisioned a kind of Google for robots. For humans, the web allows a couple of major types of knowledge sharing: (1) a person provides some description of the topic of interest (generally textual), and the system provides URLs related to the topic, and (2) a person can activate an external program which is run on a local Java interpreter. For a robot this corresponds to: (1) providing some key information based on text, images, or other sensed property of the entity, and this results in some robot digestible form of related information, and (2) a robot should be able to obtain physical behavior information (e.g., how to pick up a book), probably in the form of some standard reference language which will be interpreted on the robot's body. More recently, major efforts along these lines have been initiated, the most notable being RoboEarth [46] which is described as follows:

> At its core, RoboEarth is a World Wide Web for robots: a giant network and database repository where robots can share information and learn from each other about their behavior and their environment. Bringing a new meaning to the phrase "experience is the best teacher," the goal of RoboEarth is to allow robotic systems to benefit from the experience of other robots, paving the way for rapid advances in machine cognition and behavior, and ultimately, for more subtle and sophisticated human-machine interaction.

While this is a grand enterprise that may indeed lead to sharing at the knowledge level, and may well provide access to human knowledge, it is unclear that there is an adequate semantic basis for robots to share such knowledge. There is evidence that even sharing standard object representation models across robot platforms is difficult [18]. Moreover, great reliance is placed on human programmers to provide ontologies, as well as the frameworks for any form of sharing (e.g., sensor models, maps, coordinate frames, etc.).

Another major issue for robot knowledge sharing is the cognitivist vs. dynamical systems divide. (See Vernon et al. [44] for a survey of cognitive architectures for robots.) Cognitivist robot architectures determine actions based on syntactic manipulation of symbol tokens derived from perception; adaptation is the acquisition of knowledge; motivation is derived from some impasse to be resolved, and inter-agent actions depend on the ontology. Dynamical systems architectures, on the other hand, are some form of concurrent self-organizing network with global system states which construct skills in response to (or to cause) perturbations; motivation consists of expanding the interaction space, and inter-agent actions depend on the embodiment of the robot. Each approach has its own pros and cons (see the survey!), but here the issue is that it is desirable that robots of all cognitive types be able to share knowledge with each other and with humans. Moreover, most robots may eventually be some mix of these, with dynamical systems at the lower control levels, and symbolic representations at the higher planning levels – knowledge sharing in some form is needed across these levels as well.

If robots are to share knowledge then, it is essential to understand how knowledge is acquired and represented by robots. One approach is to simply provide some general learning mechanisms and let the robot gather sensorimotor data to discover the world. Alternatively, some innate knowledge may be embedded in the robot and this provides a description of all the entities of interest to the robot as well as how to interact with them (e.g., generate control actions) based on sensor input. Most systems lie somewhere in between these two extremes. Another major decision is whether or not to provide each robot with a description of itself or allow it to solve the sensorimotor reconstruction problem to achieve this. Finally, if new knowledge can be acquired by the robot, then the nature of this learning and its representation impacts the possibility of sharing. That is, physical symbol systems can share automata, formulas, etc., while artificial neural networks can share topology and weights, and dynamical systems can share equations and other forms of information (e.g., dominant frequencies, phase and gain values, etc. from which phase or state space symmetries can be detected and exploited).

This then leads to some of the key questions in robot knowledge sharing: (1) How much knowledge, including how to acquire knowledge, is innate (i.e., provided when built), and (2) what should this innate knowledge be? Of course, Plato held that all knowledge was innate and based on transcendent forms (invariant and unitary objects which describe the invariant relations that constitute individual objects). More recently, Chomsky [7], Pinker [36] and others have proposed that various aspects of human cognitive ability are innate (provided for genetically). Fortunately, we do not need to solve the nativist vs. empiricist debate as it applies to humans. Moreover, almost all robot cognitive architectures include some innate knowledge: subsumption architectures have innate knowledge in the form of their hardware: RoboEarth, Armar III, Soar, etc. have innate knowledge about sensors, actuators, object models, planning, language, etc. Our view is that it is too cumbersome for humans to provide all the necessary knowledge for a robot to perform robustly in the world; however, we do believe that the notion of invariance is key to providing a grounded semantics for robot knowledge. Note that there have been some recent moves to include low-level continuous representations in the more standard symbolic cognitive architectures (see Laird [24, 50] who adds a continuous representation to Soar, and also Choi [6] who propose ICARUS, a symbolic cognitive architecture for a humanoid robot); however, the high-level concepts of these systems do not arise through the sensorimotor process.

3 Symmetry in Cognition

Symmetry [49] plays a deep role in our understanding of the world in that it addresses key issues of invariance, and as noted by Viana [45]: "Symmetry provides a set of rules with which we may describe certain regularities among experimental objects." Symmetry to us means an invariant, and by determining operators which

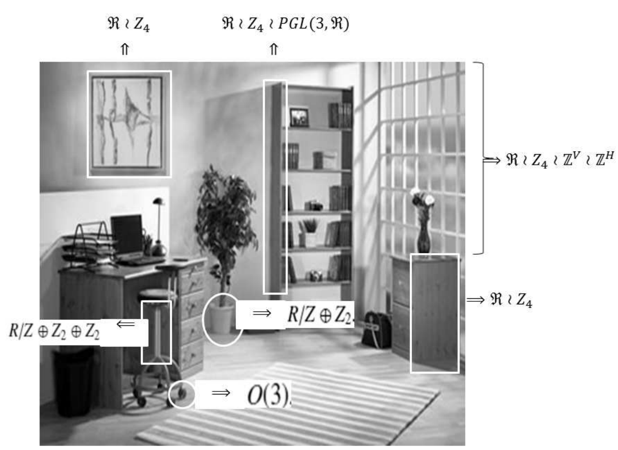

Fig. 1 Wreath Product Descriptions of an Office Scene (squares: $\Re \wr Z_4$, cone: $\Re \setminus Z \setminus \oplus Z_2$, sphere: $O(3)$, cylinder: $\Re \setminus \oplus Z_2 \oplus Z_2$, grid: $\Re \wr Z_4 \wr \mathbb{Z}^V \wr \mathbb{Z}^H$).

leave certain aspects of state invariant, it is possible to either identify similar objects or to maintain specific constraints while performing other operations (e.g., move forward while maintaining a constant distance from a wall). For an excellent introduction to symmetry in physics, see [8]. We have shown how to use symmetry in range data analysis for grasping [10]. Popplestone and Liu showed the value of this approach in assembly planning [27], while Selig has provided a geometric basis for many aspects of advanced robotics using Lie algebras [38, 39]. Recently, Popplestone and Grupen [37] gave a formal description of general transfer functions (GTF's) and their symmetries. In computer vision, Michael Leyton has described the exploitation of symmetry [25] and the use of group theory as a basis for cognition [26], and we expand on his approach here.

Leyton [26] argues that the *wreath group product* is a basic representation for cognition as stated in the *Grouping Principle*: "Any perceptual organization is structured as an n-fold wreath product $G_1 \wr \ldots \wr G_n$" and argues that "human perceptual and motor systems are both structured as wreath products." For a clear introduction to wreath products see [5] or [29]; the latter defines the wreath product as:

> Given a group $\mathscr{G}$ and a permutation group $\mathscr{H} \subseteq S_n$ (the symmetric group on n objects), the *wreath product* $\mathscr{G} \wr \mathscr{H} = \{(g,h) \mid g \in \mathscr{G}^n, h \in \mathscr{H}\}$ is a group under the operation defined by:
> $$(g',h') = (a_1,\ldots,a_n,k)(b_1,\ldots,b_n,m) := (a_1 b_{(m)^{-1}(1)},\ldots,a_n b_{(m)^{-1}(n)},km)$$

(Of course, the unrestricted wreath product may also be used.) The wreath product essentially represents all possible group actions of $\mathscr{H}$ on $\mathscr{G}$. Figure 1 depicts our goal of producing wreath product descriptions from a scene. In general, such structural descriptions can be recovered from 1-D, 2-D and 3-D data. (Note that in this figure: Z^V denotes a translation in the vertical axis, Z^H denotes a translation in the horizontal axis, Z_4 is the cyclic group of 4 elements, *PGL* is the general linear group, and the R/Z and Z_2 products help describe how to produce truncated cylinders and cones. The main point is that sensor signal data can be reduced to short expressions as geometric symmetries.)

Operationally, our hypothesis is that *group theoretic representations (G-Reps)* inform cognitive activity; *wreath products* as suggested by Leyton [25, 26] are a key part of *G-reps*. We describe here symmetry-based signal analysis and concept formation in sensorimotor reconstruction based on 1-D signals. A schematic view of our proposed symmetry-based affordance architecture (the *Symmetry Engine*) is given in Figure 2. The successful demonstration of this approach will constitute a major advance in the field of cognitive autonomous agents, and will also motivate joint research programs into human cognition.

4 Sensorimotor Reconstruction

As pointed out by Weng [48], a major research question in autonomous mental development is "how a system develops mental capabilities through autonomous real-time interactions with its environment by using its sensors and effectors (controlled by an intrinsic development program coded in the genes or designed in by hand)." Thus, a representation is sought derived from sensorimotor signals as well as the grouping of such signals as processing takes place. Note that this assumes that no

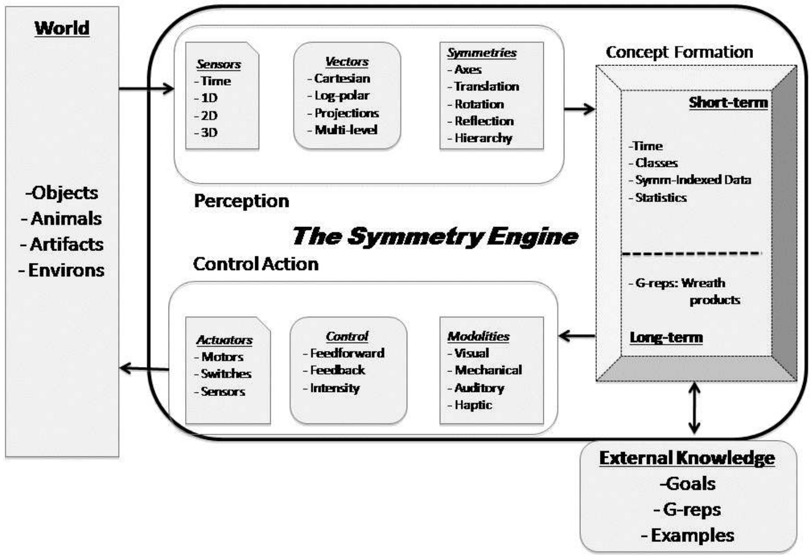

Fig. 2 The Symmetry Engine. *Perception* requires an appropriate set of operators to construct *G-reps*; this includes vector constructors, symmetry detectors, and symmetry-based data indexing and variance operators. *Control action* requires the ability to map *G-reps* onto action sequences to achieve desired results in the world. *Concept Formation* operators allow the exchange of *G-reps* with other agents. Finally, the *Human Machine Interface* (HMI) will exploit human symmetry perception, as well as *G-rep* properties to achieve context-aware integrative display of information.

coordinate frames exist in this setting; see [41] for a discussion of coordinate frames in biological systems. Asada et al. [3] give a good account of the development of body representations in biological systems and maintain that ”motions deeply participate in the developmental process of sensing and perception.” They review data ranging from spinal reflexes with fixed motor patterns, to motion assembly, to mixed motion combinations in the cerebrum. Lungarella [28] also has much to say on this issue, and of great interest here, states that ”spontaneous activity in newborns are not mere random movements ... instead organized kicks, arm movements, short phase lags between joints ... may induce correlations between sensing and motor neurons.”

Early on, Pierce [35] described an approach to learning a model of the sensor set of an autonomous agent. Features are defined in terms of raw sensor data, and feature operators are defined which map features to features. The goal is to construct a perceptual system for this structure. One of the fundamental feature operators is the *grouping operator* which assigns features to a group if they are similar. This work was extended to spatio-visual exploration in a series of papers [30, 31, 35]. For a detailed critique of Pierce’s work, see [11]. Olsson extended this work in a number of papers [19, 20, 21, 22, 23, 32, 33]. He used information theoretic measures for sensorimotor reconstruction, and no innate knowledge of physical phenomena or the sensors is assumed. Like Pierce, Olsson uses random movements to build the representation and learn the effect of actions on sensors to perform visually guided movements. The major contributions are the analysis of information theoretic measures and motion flow. O'Regan and Noë [34] use the term *sensorimotor contingencies* and give an algorithm which can determine the dimension of the space of the environment by "analyzing the laws that link motor outputs to sensor inputs"; their mathematical formulation is elegant.

4.1 Symmetry Detection in 1-D Signals

A symmetry defines an invariant. The simplest invariant is identity. This can apply to an individual item, i.e., a thing is itself, or to a set of similar objects. In general, an invariant is defined by a transformation under which one object is mapped to another. Sensorimotor reconstruction can be more effectively achieved by finding such symmetry operators on the sensor and actuator data (see also [4, 17]).

Invariants are very useful things to recognize, and we propose that various types of invariant operators provide a basis for cognitive functions, and that it is also useful to have processes that attempt to discover invariance relations among sensorimotor data and subsequently processed versions of that data.

Assume a set of sensors, $\mathscr{S} = \{S_i, i = 1 \ldots n_{\mathscr{S}}\}$ each of which produces a finite sequence of indexed sense data values, s_{ij} where i gives the sensor index and j gives an ordinal temporal index, and a set of actuators, $\mathscr{A} = \{A_i, i = 1 \ldots n_{\mathscr{A}}\}$ each of which has a finite length associated control signal, A_{ij}, where i is the actuator index and j is a temporal ordinal index of the control values.

Here we are interested in determining the similarity of sensorimotor signals. Thus, the type of each sensor as well as the relation to motor control actions play a role. It is quite possible that knowledge of the physical phenomenon that stimulates a sensor may also be exploited to help determine the structure of the sensor system and its relation to motor action and the environment [12].

We suppose that certain 1D signal class structures are important and are known a priori to the agent (i.e., that there are processes for identifying signals of these types). Given an isometry that maps a 1-D signal onto itself, there are 3 possibilities:

1. All points map to themselves (the *identity* symmetry).
2. Only one point maps to itself (the *reflection* symmetry). E.g., this is the case for the histogram of a Gaussian sample.
3. No point maps to itself (a *translation* symmetry). A translation can be either continuous (a linear signal) or discrete (a periodic signal).

We have developed algorithms to detect these symmetries and use them to classify sensor and actuator types in the sensorimotor reconstruction problem (see [16]). This allows sensor classification without any actuation (i.e., much lower energy expenditure), and achieves much greater classification correctness compared to previous methods.

The output of the 1-D symmetry analysis is one of:

- Z_1: a non-symmetric base shape b.
- Z_2: a basic shape b with reflexive symmetry.
- $\Re$: a continuous translational signal; i.e., a line with slope m and intercept b.
- $\mathbb{Z}$: a periodic signal with base shape b and period T.

Note that symmetry analysis may be applied to transformed signals (e.g., to the histogram of a signal; a Gaussian sample should result in the structural type Z_2).

Thus, a first level symmetry is one that characterizes a single signal as belonging to one of these categories. Of course, composite signals can be constructed from these as well.

Next, pairwise signal symmetries can exist between signals in the same class, and if so, they will be grouped:

- *linear (y = ax + b)*
  - same line: $a_1 = a_2$, $b_1 = b_2$
  - parallel: $a_1 = a_2$, $b_1 \neq b_2$
  - intersect in point: rotation symmetry about intersection point
- *periodic (y(t+T) = y(t))*
  - same period
  - same Fourier coefficients
- *Gaussian* $(N(\mu,\sigma^2))$
  - same mean
  - same variance

4.2 The Sensorimotor Reconstruction Process

The sensorimotor reconstruction process consists of the following steps: (1) perform actuation command sequences, (2) record sensor data, (3) determine sensor equivalence classes, and (4) determine sensor-actuator relations. An additional criterion is to make this process as efficient as possible.

Olsson, Pierce and others produce sensor data by applying random values to the actuators for some preset amount of time, and record the sensor sequences, and then look for similarities in those sequences. This has several problems: (1) there is no guarantee that random movements will result in sensor data that characterizes similar sensors, (2) there is no known (predictable) relation between the actuation sequence and the sensor values, and (3) the simultaneous actuation of multiple actuators confuses the relationship between them and the sensors.

To better understand sensorimotor affects, a systems approach is helpful. That is, rather than giving random control sequences and trying to decipher what happens, it is more effective to hypothesize what the actuator is (given limited choices) and then provide control inputs for which the effects are known. Such hypotheses can be tested as part of the developmental process. The basic types of control that can be applied include: none, impulse, constant, step, linear, periodic, or other (e.g., random).

Next, consider sensors. Some may be time-dependent (e.g., energy level), while others may depend on the environment (e.g., range sensors). Thus, it may be possible to classify ideal (noiseless) sensors into time-dependent and time-independent by applying no actuation and looking to see which sensor signals are not constant (this assumes the spatial environment does not change). This also applies to noisy sensors in that it may be more useful to not actuate the system, and then classify sensors based on their variance properties. That is, in realistic (with noise) scenarios, it may be possible to group sensors without applying actuation at all.

Consider Pierce's sensorimotor reconstruction process. If realistic noise models are included, the four types of sensors in his experiments (range, broken range, bearing and energy) can all be correctly grouped with no motion at all. (This assumes some energy loss occurs to run the sensors.) All this can be determined just using the 1-D symmetries described above and the means and variances of the sensor data sequences. This leads to the following algorithms:

Algorithm SBSG: Symmetry-Based Sensor Grouping

```
1. Collect sensor data for given period
2. Classify Sensors as Basic Types
3. For all linear sensors
   a. Group if similar regression error
4. For all periodic sensors
  a. Group if similar Period
5. For all Gaussian sensors
  a. Group if similar variance
```

This algorithm assumes that sensors have an associated noise. Note that this requires no actuation and assumes the environment does not change. Finally, the similarity test for the above algorithm depends on the agent embodiment. Note that in 4a and 5a above, a similarity measure must be established; this depends on the particular application and needs.

Algorithm SBSR: Symmetry-Based Sensorimotor Reconstruction

```
1. Run single actuator and
   collect sensor data for given period
2. For each set of sensors of same type
   a. For each pair
      i. If translation symmetry holds
         Determine shift value
         (in actuation units)
```

This determines the relative distance (in actuation units) between sensors. E.g., for a set of equi-spaced range sensors, this is the angular offset.

4.3 Comparison to Pierce's Work

A set of simulation experiments are described in Chapter 4 of Pierce's dissertation [35]. The first involves a mobile agent with a set of range sensors, a power level sensor, and four compass sensors. The sensors are grouped and then a structural layout in 2D is determined. The second experiment concerns an array of photoreceptors. Here we examine the first experiment, and in particular, the group generator.

The basic setup involves a $6 \times 4\ m^2$ rectangular environment with a mobile robot defined as a point. The robot is equipped with 29 sensors all of which take values in the range from zero to one. Sensors 1 to 24 are range sensors which are arranged in an equi-spaced circle aiming outward from the robot. Range sensor 21 is defective and always returns the value 0.2. Sensor 25 gives the voltage level of the battery while sensors 26 to 29 give current compass headings for East, North, West and South, respectively. The value is 1 for the compass direction nearest the current heading and zero for the other compass sensors. There are two motors, a_0 and a_1, to drive the robot, and these can produce a maximum forward speed of 0.25 m/sec, and a maximum rotation speed of 100 degrees/sec. We assume that the values of the motors range from -1 to 1, where -1 produces a backward motion and 1 produces a forward motion (more specifically, assume the rotational axis of the tracks is aligned with the y-axis; then a positive rotation moves z into x and corresponds to a positive rotation about y in the coordinate frame).

Some details of the motion model are left unspecified; therefore we use the following model:

```
if a0>= 0 and a1>=0
then robot moves forward min(a0,a1)*0.25 m/sec
     robot rotates ((a0-a1)/2)*100 degrees/sec

elseif a0<=0 and a1<=0
then robot moves backward abs(max(a0,a1))*0.25 m/sec
     robot rotates ((a0-a1)/2)*100 degrees/sec

elseif a0>0 and a1<0
then robot rotates ((a0-a1)/2)*100 degrees/sec

end
```

Moreover, if the robot attempts to move out of the rectangular environment, no translation occurs, but rotation does take place.

Two pairwise metrics are defined (vector and PDF distances), and based on these the sensors are grouped pairwise. Then the transitive closure is taken on these. Pierce runs the simulation for 5 simulated minutes and reports results on the sample data generated from that run. Based on the samples generated from this run, the group generator produces seven groups:

```
Range: {1,2,3,4,5,6,7,8,9,10,11,12,13,
        14,15,16,17,18,19,20,22,23,24}
Defective range: {21}
Battery Voltage: {25}
Compass (East):  {26}
Compass (North): {27}
Compass (West):  {28}
Compass (South): {29}
```

It is not clear why range sensors are grouped, but compass sensors are not, nor why a run of five minutes was selected. In our attempt to replicate this experiment, we tried runs of various time lengths and found that the grouping correctness rose to a maximum at about five minutes, but never got a perfect result.

4.4 Symmetry-based Grouping Operator

Any simulation experiment should carefully state the questions to be answered by the experiment and attempt to set up a valid statistical framework. In addition, the sensitivity of the answer to essential parameters needs to be examined. Pierce does not explicitly formulate a question, nor name a value to be estimated, but it seems clear that some measure of the correctness of the sensor grouping would be appropriate. From the description in the dissertation, Pierce ran the experiment once for 5 minutes of simulated time, and obtained a perfect grouping solution.

From this we infer that the question to be answered is:

Grouping Correctness: What is the correctness performance of the proposed grouping generator?

This requires a definition of correctness for performance and we propose the following:

Correctness Measure: Given (1) a set of sensors, $\{S_i, i = 1 : n\}$ (2) a correct grouping matrix, G, where G is an n by n binary valued matrix with $G(i, j) = 1$ if sensors S_i and S_j are in the same group and $G(i, j) = 0$ otherwise, and (3) H an n by n binary matrix which is the result of the grouping generator, then the grouping correctness measure is:

$$\mu_G(G,H) = \sum_{i=1}^{n}\sum_{j=1}^{n}[(\delta_{i,j})/n^2]$$

$$\delta_{i,j} = 1\text{if G(i,j)==H(i,j); 0 otherwise}$$

Note that we define G here for the purpose of evaluation of the method, but a robot agent will need to validate any groupings that it discovers. This involves some form of learning process, and we do not deal with that here. However, we believe that this can be based on how well affordances work which depend on the grouping.

4.4.1 Sensor Grouping with Noise (No actuation)

Assume that the sensors each have a statistical noise model. The real-valued range sensors have Gaussian noise sampled from a $\mathcal{N}(0,1)$ distribution (i.e., $v_{sample} = v_{true} + \omega$. The binary-valued bearing sensors have salt and pepper noise where the correct value is flipped p% of the time. Finally, the energy sensor has Gaussian noise also sampled from $\mathcal{N}(0,1)$. (The broken range sensor returns a constant value.)

Based on this, the grouping correctness results using SBSG are given in Figure 3. Sensor data sampling time was varied from 1 to 20 seconds for binary noise of 5%, 10% and 25%, and Gaussian variance values of 0.1, 1, and 10. Ten trials were run for each case and the means are shown in the figure. As can be seen, perfect sensor grouping is achieved after 20 seconds without any actuation cost. Previous methods required driving both wheels for a longer time and they cost about $30k_{a/s}$ more in energy than our method ($k_{a/s}$ is the actuation to sensing cost ratio).

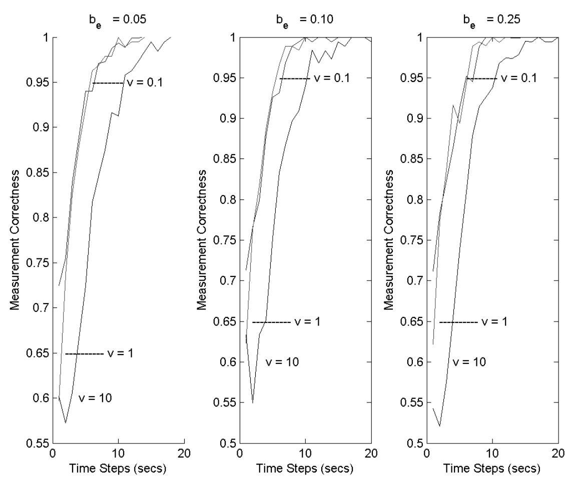

Fig. 3 Grouping Correctness vs. Number of Samples for SBSG; left to right are for binary salt and pepper noise of 5%, 10%, and 25%; curves for 0.1, 1.0, and 10.0 variance are given in each plot.

4.4.2 Sensor Grouping (Actuated)

Given a set of sensors that characterize the group operation nature of an actuator (in this case rotation), the sensors can be grouped based on the fact that similar sensors produce data that has a translation symmetry along the temporal axis. Figure 4 shows representative data for the simulated range and compass sensors. The simple determination of a translation symmetry between signals allows both grouping (i.e., the signals match well at some time offset), and the angular difference between the sensors (given by the t_{offset} at which the symmetry occurs); t_{offset} is proportional to the angle between the the sensors in terms of actuation units. Figure 5 shows the perfect grouping result with noise of 1% in the compass sensor data and 0.1 variance in the range sensor data (the figure shows a 29x29 similarity matrix where white indicates sensors are in same group, and black indicates that are not).

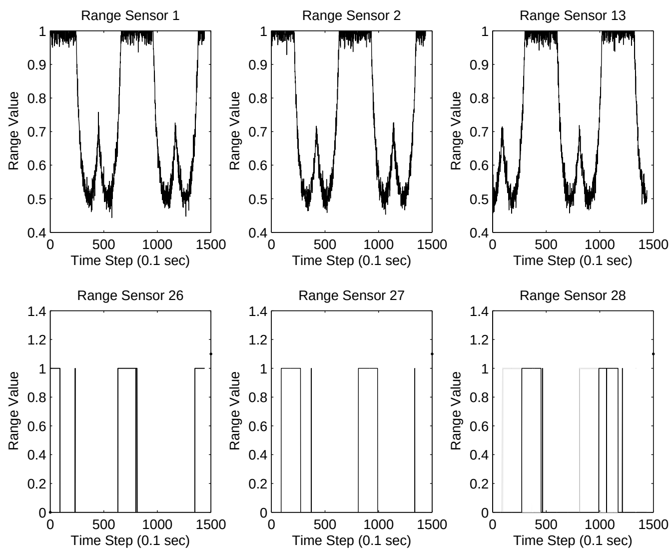

Fig. 4 Sensor data showing translation symmetry: Row 1 shows sensors 1, 2, and 13; Row 2 shows compass sensors 27,28, and 29.

4.4.3 Unactuated Physical Experiment

We have performed experiments with physical sensors to validate the proposed approach. Data was taken for both the static case (no actuation) and the actuated case (camera rotation). Two sensors were used in this experiment: a camera and a microphone. The camera was set up in an office and a sequence of 200 images was taken at a 10Hz rate. Figure 6 shows one of these images. The 25x25 center set of pixels from the image comprise a set of 625 pixel signals each of length 200. An example trace and its histogram are given in Figure 7. As can be seen, this is qualitatively a Gaussian sample. Figure 8 shows a 200 sequence signal of microphone data, and its histogram which also looks Gaussian.

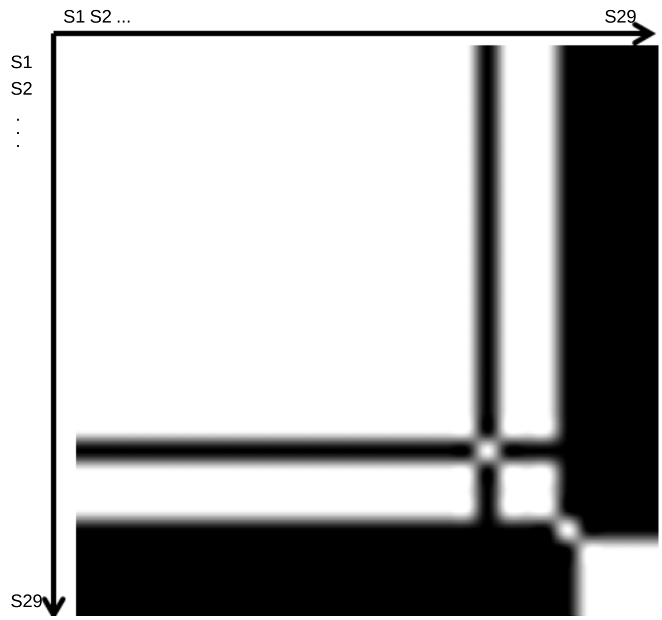

Fig. 5 Grouping Matrix: 29×29 binary matrix; sensors 1-24 are range sensors (sensor 21 returns constant value); 25 is energy; 26-29 are compass sensors.

The application of our symmetry detectors classified all pixel and microphone signals as Gaussian signals, and grouped the pixel signals separately from the microphone due to the difference in their variance properties.

4.4.4 Actuated Physical Experiment

Of course, actuation can help understand the structure of the sensorimotor system. For example, consider what can be determined by simply rotating a two-wheeled robot that has a set of 22 range sensors arranged equi-spaced on a circle. Assume that the control signal results in a slow rotation parallel to the plane of robot motion (i.e., each range sensor moves through a small angle to produce its next sample) and rotates more than 2π radians. Then each range sensor produces a data sequence that is a shifted version of each of the others – i.e., there is a translation symmetry (of periodic signals) between each pair. The general problem is then:

General Symmetry Transform Discovery Problem: Given two sensors, S_1 and S_2, with data sequences T_1 and T_2, find a symmetry operator σ such that $T_2 = \sigma(T_1)$.

We also took a set of images by rotating the camera by five degree increments for 720 degrees (see Figure 9 for the first eight of the 128 images in the rotated sequence). Domain translation symmetry allows the identification of all the pixel

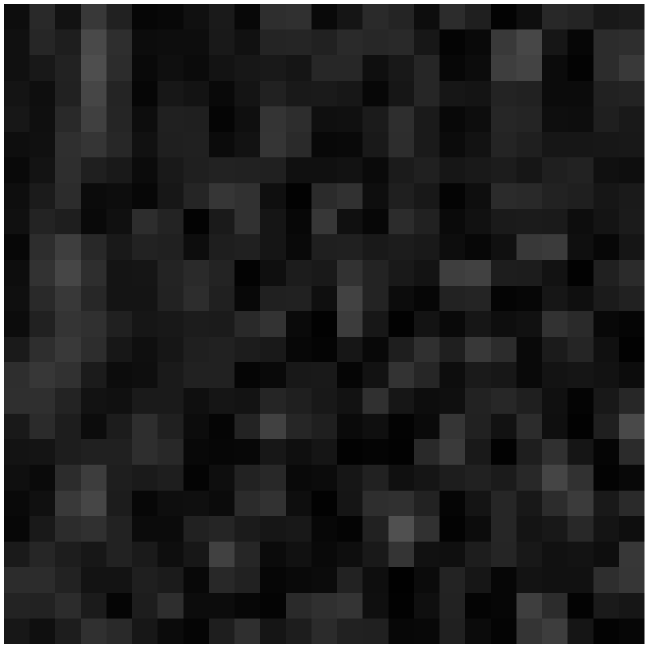

Fig. 6 The 25x25 Center Pixels from One of the 200 Static Images.

signals along a row as similar to each other (i.e., they are all in the plane of the rotation). Due to the translation amount, the offset between the signals is also discovered. Moreover, due to the fact that individual signals are classified as periodic (with period 64), it is determined that the actuator is performing a rotation about the axis orthogonal to the camera's optical axis (Figure 10 shows the overlay of the two periodic 64-element pieces of the 1-D signal for the center pixel of the sequence of 128 images).

The groupings found here are mainly useful to allow the discovery of, e.g., the pixels in a specific camera, or the rangels in a range finder, etc., and nothing guarantees that distinct range finders will be grouped. For example, if their noise characteristics are dissimilar (based on the selected similarity measure), then they will not be grouped. On the other hand, groupings should be pretty consistent even with changes in environmental conditions since the groupings are mainly based on structure and noise properties. Finally, we have not yet considered combined translation and rotation, but run actuators independently. This is an interesting topic for future research.

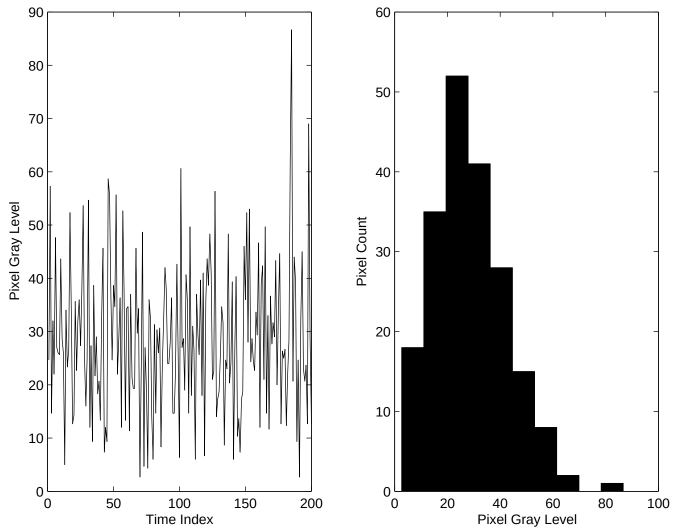

Fig. 7 Trace and Histogram of the 200 Pixel Values of the Center Pixel of the Images.

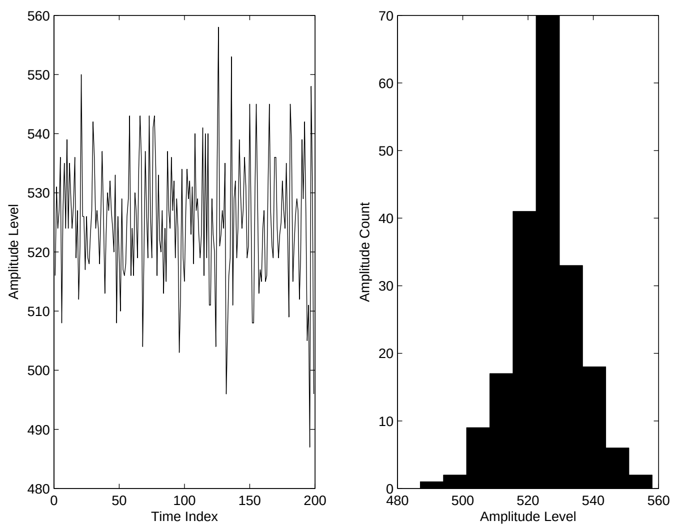

Fig. 8 Trace and Histogram of the 200 Amplitude Values of the Microphone Data.

5 Conclusions and Future Work

We propose symmetry theory as a basis for sensorimotor reconstruction in embodied cognitive agents and have shown that this allows the identification of structure with simple and elegant algorithms which are very efficient. The exploitation of noise structure in the sensors allows unactuated grouping of the sensors, and a simple

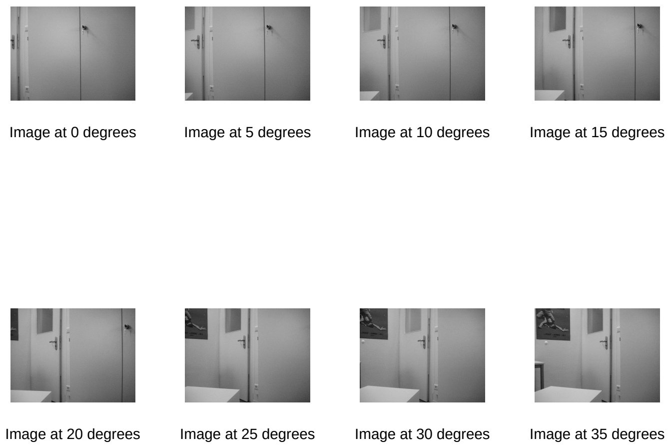

Fig. 9 The First Eight Images in the 128 Image Sequence over 720 Degree Rotation

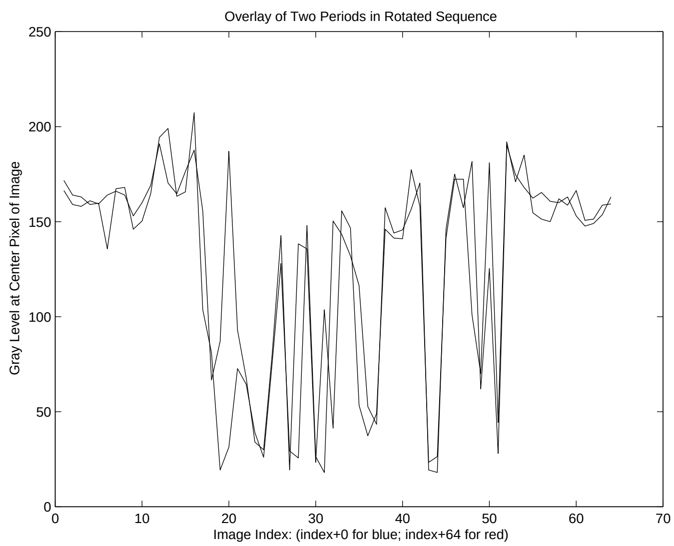

Fig. 10 Overlay of the Two Recovered Periodic Parts of the Two Revolutions of Image Data.

one actuator rotation permits the recovery of the spatial arrangement of the sensors. This method was shown to hold for physical sensors as well. This form of structural bootstrapping involves both the identification of instances of structural prototypes (i.e., specific 1-D symmetries), as well as the subsequent classification of a broader type of entity (e.g., range sensors).

Several directions remain to be explored:

1. Consider rotational actuators; these can be seen to define a group in the following way: any specific rotation is an element of the group set, and application of

rotation is the operator. Group properties can be seen to hold in that (i) the sequential application of two rotations is a rotation, (ii) the opposite rotation is the inverse element, (iii) the application of no actuation is the identity element, and (iv) associativity holds. [Note that rotation in just one sense forms a group, and various combinations of actuators may form larger groups - e.g., two wheels.]
$\rightarrow$ The analysis of actuators as specific group operators requires study.
2. Higher-dimensional symmetries offer many opportunities for research. For example, the transformation from spatial image layout to log-polar form allows 1D symmetries to be sought which characterize object scaling and rotation.
$\rightarrow$ The analysis of higher-dimensional symmetries requires study.
3. Higher-level sensorimotor symmetries will allow the conceptualization of physical objects in terms of sensorimotor sequences characterized by some invariant (e.g., stand-off distance in circumlocuting the object).
$\rightarrow$ The analysis of symmetries in sensorimotor interactions with the environment requires study.
4. Finally, we are instrumenting a set of mobile robots with range and other sensors and a series of experiments will be conducted to study these broader issues.
$\rightarrow$ Experimental studies in broader environmental interaction are required.

Acknowledgments
This material is based upon work supported by the National Science Foundation under Grant No. 1021038, and is also based in part on [13].

References

1. A. Chella and M. Cossentino and R. Pirrone and A. Ruisi. Modeling Ontologies for Robotic Environments. In *Proceedings of the Fourteenth International Conference on Software Engineering and Knowledge Engineering*, pages 15–19, 2002.
2. F. Amigoni and M.A. Neri. An Application of Ontology Technologies to Robotic Agents. In *Proceedings of the IEEE/WIC/ACM International Conference on Intelligent Agent Technology*, IAT '05, pages 751–754, Washington, DC, USA, 2005. IEEE Computer Society.
3. M. Asada, K. Hosoda, Y. Kuniyoshi, H. Ishiguro, T. Inui, Y. Yoshikawa, M. Ogino, and C. Yoshida. Cognitive Developmental Robotics: A Survey. *IEEE Transactions on Autonomous Mental Development*, 1(1):12–34, 2009.
4. F. Bullo and R. M. Murray. Proportional Derivative (PD) Control On The Euclidean Group. In *Proceeding of the European Control Conference*, pages 1091–1097, 1995.
5. W.Y. Chang. Image Processing with Wreath Products. Master's thesis, Harvey Mudd College, Claremont, CA, May 2004.
6. D. Choi, Y. Kang, H. Lim, and B.-J. You. Knowledge-based Control of a Humanoid Robot. In *Proceedings of the IEEE International Conference on Intelligent Robots and Systems*, St. Louis, MO, September 2009.
7. N. Chomsky. Biolinguistic Explorations: Design, Development, Evolution. *International Journal of Philosophical Studies*, 15(1):1–21, March 2007.
8. J.P. Elliott and P.G. Dawber. *Symmetry in Physics*. The Macmillan Press, London, UK, 1979.
9. X. Fan. RobotShare: A Google for Robots. Master's thesis, University of Utah, Salt Lake City, UT, May 2008.
10. R. Grupen and T.C. Henderson. Apparent Symmetries in Range Data. *Pattern Recognition Letters*, 7:107–111, 1988.

11. T. Henderson and H. Peng. A study of pierce's group generator. Technical Report UUCS-10-001, The University of Utah, December 1984.
12. T.C. Henderson. *Computational Sensor Networks*. Springer Verlag, New York, NY, 2009.
13. T.C. Henderson, E. Cohen, A. Joshi, E. Grant, M. Draelos, and N. Deshpande. Symmetry as a Basis for Perceptual Fusion. In *Proceedings of the 2012 IEEE International Conference on Multisensor Fusion and Integration*, Hamburg, Germany, September 2012.
14. T.C. Henderson and X. Fan. Robot Semantic Web. In *IEEE IROS Workshop on Robot Semantic Web, IROS 2007*, San Diego, CA, 2007. IEEE.
15. T.C. Henderson and X. Fan. RobotShare: A Google for Robots. *Special Issue on Cognitive Humanoid Robots of the International Journal of Humanoid Robotics*, 5:311–329, June 2008.
16. T.C. Henderson, H. Peng, C. Sikorski, N. Deshpande, and E. Grant. Symmetry: A Basis for Sensorimotor Reconstruction. Technical Report UUCS-11-011, The University of Utah, May 2011.
17. J.C. Kinsey and L.L. Whitcomb. Adaptive Identification on the Group of Rigid Body Rotations. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 3256–3261. IEEE Press, 2005.
18. Z. Kira. *Communication and Alignment of Grounded Symbolic Knowledge among Heterogeneous Robots*. PhD thesis, Georgia Tech, Atlanta, GA, 2010.
19. C.L. Nehaniv L. Olsson and D. Polani. Sensory Channel Grouping and Structure from Uninterpreted Sensor Data. In *Proceedings of NASA/DoD Conference on Evolvable Hardware*, pages 153–160, Seattle, WA, 2004. IEEE Computer Society Press.
20. C.L. Nehaniv L. Olsson and D. Polani. The Effects on Visual Information in a Robot in Environments with Oriented Contours. In L. Berthouze, H. Kozima, C.G. Prince, G. Sandini, G. Stojanov, and G. Metta, editors, *Proceedings of the Fourth International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, pages 83–88, Genoa, Italy, 2004. Lund University Cognitive Studies.
21. C.L. Nehaniv L. Olsson and D. Polani. Discovering Motion Flow by Temporal-Informational Correlations in Sensors. In L. Berthouze, H. Kozima, C.G. Prince, G. Sandini, G. Stojanov, G. Metta, and C. Balkenius, editors, *Proceedings of the Fifth International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, pages 117–120, Nara, Japan, 2005. Lund University Cognitive Studies.
22. C.L. Nehaniv L. Olsson and D. Polani. From Unknown Sensors and Actuators to Visually Guided Movement. In *Proceedings of the International Conference on Development and Learning*, pages 1–6, Osaka, Japan, 2005. IEEE Computer Society Press.
23. C.L. Nehaniv L. Olsson and D. Polani. Sensor Adaptation and Development in Robots by Entropy Maximization of Sensory Data. In *Proceedings of the International Symposium on Computational Intelligence in Robotics and Automation*, pages 587–592, Espoo, Finland, 2005. IEEE Computer Society Press.
24. J. Laird. Extending the Soar Cognitive Architecture. In *Proceedings of the Artificial General Intelligence Conference*, Memphis, TN, March 2008.
25. M. Leyton. *Symmetry, Causality, Mind*. MIT Press, Cambridge, MA, 1992.
26. M. Leyton. *A Generative Theory of Shape*. Springer, Berlin, 2001.
27. Y. Liu and R. Popplestone. A Group Theoretic Formalization of Surface Contact. *International Journal of Robotics Research*, 13(2):148–161, 1994.
28. M. Lungarella, G. Metta, R. Pfeifer, and G. Sandini. Developmental Robotics: A Survey. *Connection Science*, 15(4):151–190, 2003.
29. P. McMahan. Wreath-Product Polytopes. Master's thesis, Reed College, Portland, OR, May 2005.
30. J. Modayil and B. Kuipers. Autonomous Development of a Grounded Object Ontology by a Learning Robot. In *Proceedings of the Twenty-Second Conference on Artificial Intelligence*, Vancouver, Canada, 2007. AAAI.
31. J. Modayil and B. Kuipers. The Initial Development of Object Knowledge by a Learning Robot. *Robotics and Autonomous Systems*, 56(11):879–890, 2008.

32. L. Olsson, C.L. Nehaniv, and D. Polani. Measuring Informational Distances between Sensors and Sensor Integration. Technical report 431, University of Hertfordshire, Hertfordshire, England, 2005.
33. L. Olsson, C.L. Nehaniv, and D. Polani. From Unknown Sensors and Actuators to Actions Grounded in Sensorimotor Perceptions. *Connection Science*, 18:121–144, 2006.
34. J.K. O'Regan and A. Noë. A Sensorimotor Account of Vision and Visual Consciousness. *Behavioral and Brain Sciences*, 24:939–1031, 2001.
35. D.M. Pierce. *Map Learning with Uninterpreted Sensors and Effectors*. PhD thesis, The University of Texas at Austin, Austin, TX, May 1995.
36. S. Pinker. *The Language Instinct*. Harper Collins, NY, NY, 1994.
37. R. Popplestone and R. Grupen. Symmetries in World Geometry and Adaptive Behaviour. In *Proceedings of the Workshop on Algebraic Frames for the Perception Action Cycle*, pages 269–283, Kiel, Germany, 2000.
38. J.M. Selig. Lie Groups and Lie Algebras in Robotics. In J. Byrnes, editor, *Proceedings of the NATO Advanced Study Institute on Computational Noncommutative Algebra and Applications*, Il Ciocco, Italy, 2004. Kluwer.
39. J.M. Selig. *Geometric Fundamentals of Robotics*. Springer, Berlin, 2005.
40. Jivko Sinapov and Alexander Stoytchev. Toward Autonomous Learning of an Ontology of Tool Affordances by a Robot. In *Proceedings of the 23rd National Conference on Artificial intelligence - Volume 3*, pages 1828–1829. AAAI Press, 2008.
41. J.F. Soechting and M. Flanders. Moving in Three-Dimensional Space: Frames of Reference, Vectors, and Coordinate Systems. *Annual Reviews on Neuroscience*, 15:167–191, 1992.
42. I.H. Suh, G.H. Lim, W. Hwang, H. Suh, J.-H. Choi, and Y.-T. Park. Ontology-based Multi-layered Robot Knowledge Framework (OMRKF) for Robot Intelligence. In *IEEE International Conference on Intelligent Robots and Systems*, San Diego, CA, 2007. IEEE.
43. M. Tenorth and M. Beetz. KNOWROB – Knowledge Processing for Autonomous Personal Robots. In *IEEE International Conference on Intelligent Robots and Systems*, St. Louis, MO, 2009. IEEE.
44. 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*, 11(2):151–180, 2008.
45. M.A.G. Viana. *Symmetry Studies*. Cambridge University Press, Cambridge, UK, 2008.
46. M. Waibel, M. Beetz, J. Civera, R. D'Andrea, J. Elfring, D. Galvez-Lopez, K. Haussermann, R. Jannsen, J.M.M. Montiel, A. Perzylo, B. Schiessle, M. Tenorth, O. Zweigle, and R. Molengraft. RoboEarth – A World Wide Web for Robots. *Robotics and Automation Magazine*, 18(2):69–82, June 2011.
47. E. Wang, Y.S. Kim, H.S. Kim, J. H. Son, S. Lee, and I.H. Suh. Ontology Modeling and Storage System for Robot Context Understanding. In R. Khosla, editor, *Proceedings of the 3rd Conference on Knowledge-Based Intelligent Information and Engineering Systems*, pages 922–929, Berlin, 2005. Springer.
48. J. Weng and I. Stockman. Autonomous Mental Development: Workshop on Development and Learning. *AI Magazine*, 23(2), 2002.
49. H. Weyl. *Symmetry*. Princeton University Press, Princeton, NJ, 1952.
50. J. Xu. New SVS Implementation. In *Soar Workshop*, 2011.