From Unknown Sensors and Actuators to Actions Grounded in Sensorimotor Perceptions

Lars Olsson[†], Chrystopher L. Nehaniv[†‡], Daniel Polani[†‡]
Adaptive Systems[†] and Algorithms[‡] Research Groups
School of Computer Science,University of Hertfordshire
College Lane, Hatfield Herts AL10 9AB, United Kingdom
{L.A.Olsson, C.L.Nehaniv, D.Polani}@herts.ac.uk

Abstract

This article describes a developmental system based on information theory implemented on a real robot that learns a model of its own sensory and actuator apparatus. There is no innate knowledge regarding the modalities or representation of the sensory input and the actuators, and the system relies on generic properties of the robot's world such as piecewise smooth effects of movement on sensory changes. The robot develops the model of its sensorimotor system by first performing random movements to create an informational map of the sensors. Using this map the robot then learns what effects the different possible actions have on the sensors. After this developmental process the robot can perform basic visually guided movement.

Keywords: developmental robotics, information theory, emergence of structure

1 Introduction

All animals display amazing capabilities and control over their own bodies. They are also often able to predict the way in which certain actions affect the environment around them. In order to achieve this, many animals go through a number of developmental stages, whereby the nervous system develops and the animal learns to control its own body and the relations between sensors and actions. In the process of learning a *sense of bodily self* in human infants, systematic exploration of the sensory and perceptual consequences of their actions plays a major role (Rochat, 1998). For example, young infants perform the same actions over and over again (Piaget, 1953), and it has been observed that newborn infants spend up to 20% of their time while awake touching their face with their hands (Korner and Kraemer, 1972). This phenomenon, in analogy with the vocal babbling in the early process of language learning, is also called *body babbling* (Meltzoff and Moore, 1997). It has also been found in experiments

with kittens that visually guided behaviour only develops when changes in visual stimulation are related to self-initiated movements performed by the kitten itself (Held and Hein, 1963). Thus, it seems important to ground actions in sensorimotor perceptions.

How to ground actions in sensorimotor perceptions in robots is studied in the areas of autonomous mental development (Weng et al., 2001) and developmental robotics, see (Lungarella et al., 2004) for a review. These fields have two major driving forces. One is to build more adaptable, autonomous, and sociable robots by drawing inspiration from developmental psychology and neuroscience. The other major driving force is to use robots as embodied tools to help the investigation of the development of neural systems and cognitive abilities. In the present article the main focus is on the former, even though the latter will also be touched upon. Methods based on information theory are presented that enable a robotic system to develop from no knowledge of its sensor modalities and arrangements, as well as actuators, to sensory-guided movement by learning relations between actions and their effect on its sensors.

The first question one needs to ask is why it is interesting desirable to let a robot learn by itself about its sensors and actuators? Traditionally, the possible actions and sensorimotor relations of robots are specified and built in to the robot by the human designer. But, as experienced in many robotic projects, it is extremely hard to design a robotic system to use its sensors in an efficient manner in challenging realistic and changing environments – even more so with more complex robots. It is also very hard as a human designer to understand the environment in which the robot is acting in from the robot's perspective. Thus, if the robot learns by itself its possible actions and their effect on its sensors, anthropomorphic bias is avoided since the actions are acquired by self-experienced sensorimotor percepts (Blank et al., 2005) in its own *Umwelt* (von Uexküll, 1956) – the sensory world of the agent. Also, a robot that acquires the skills necessary to act in the world by itself can adapt to changing conditions and embodiment since the learning of relations between actuators and sensors, the *sensorimotor laws*, can adapt over time. By developing autonomous ontogeny by grounding the actions in sensorimotor laws, the robot's repertoire is more open-ended since it is not learning a specific task, and can continue to adapt to changes and new tasks. Finally, studying how autonomous development of sensorimotor skills operates in robots can help validate or extend existing theories in human and animal development.

Building a robot capable of open-ended autonomous development is of course a complex problem and in the present article we present a unified approach to tackling a number of aspects of autonomous development. More specifically, the goal of the presented system is to develop from no knowledge of its sensors and actuators and their relations to performing structured movement guided by the discovery of informational structure in the sensorimotor system. After learning the informational structure of the sensors, the effects certain settings of the actuators have on the sensors are learned. This can be seen as learning a *forward model* (Jordan and Rumelhart, 1992) of the sensors and actuators, but via general means based on informational relations. Given a certain setting

of the actuators (which constitute an action), what effect will this action have on the sensors? For example, if the head is moved down, there will be a motion flow moving upwards in the visual sensors. This kind of learning is done autonomously by the robot, using the equivalent of body babbling in animals, here called *motor babbling*.

Motor babbling has been investigated by a number of researchers to explore the possible sensorimotor space of a specific embodiment. For example, Berthouze et al. (1998) developed a system performing basic visuo-motor behaviors which were categorized by a neural network. In Kuniyoshi et al. (2003) body babbling was used to develop a system capable of forms of imitation. Berthouze and Kuniyoshi (1998) also developed a system capable of performing unsupervised learning of visual sensorimotor patterns, where the classification used Kohonen maps. The resulting self-organized Kohonen maps showed four different categories of movement: vertical, 'in-depth motions', horizontal, and a fourth, not clearly defined, intermediate category between the horizontal and 'in-depth motions' categories. Related work by Kuperstein studied adaptive hand-eye coordination for reaching (Kuperstein, 1988). Here a neural network model was developed that learnt to reach a cylinder arbitrarily placed in space. Self-produced motor signals were used to explore many different arm positions with the cylinder in the hand and topographic sensory information stored in maps. These topographic mappings allowed the system to learn to reach only from sensory receptors and motor feedback, without a priori knowledge about object features. Learning of associations between vision and arm movements has also been studied by Andry et al. (2004). This work is similar to Kuperstein's in many respects via its use of visuo-motor maps and neural network control, but also raised the idea of co-development between sensorimotor and imitative capabilities. Stronger and Stone (2006) developed a method where a robot learns by trying different actuator settings the relationships between a sensor's values and the current state of the world, and also the relationships between the actuator settings and the rate of change of the sensor's readings.

One common feature of the methods just presented is that they all assume that the agent knows about the structure of its sensors, and that all the sensors are visual sensors or have similar properties of continuity to visual sensors. In the methods presented in the present article the focus is different since it is the agent itself that must learn the informational structure of its sensors before it can develop sensory-guided movement.

Central to the concept of discovering the informational structure and flow in sensors are the informational relationships between sensors and the statistics of the signals in the sensors. The structure of the incoming sensory signals depends on the embodiment and actions of the agent and the environment. Research into the structure of natural signals is still at an early and exploratory phase, but there are indications that signals of different sensory modalities such as acoustic waveforms, odour concentrations, and visual contrast share some statistical properties (Rieke et al., 1999). For example, in general the local contrast in natural images has the same exponential form of the probability distribution as sound pressure in musical pieces (Rieke et al., 1999). Another

commonality between signals of different modalities is coherence over time and, in many cases, spatial coherence. It was also found in Coppola et al. (1998) that man-made environments as well as natural environments have more vertical and horizontal contours than oblique angled contours.

These results all suggest that taking the statistical structure of natural signals into account can be exploited by both biological organisms as well as machines. In the early 1960s H. B. Barlow suggested (Barlow, 1961) that the visual system of animals "knows" about the structure of natural signals and uses this knowledge to represent visual signals. Laughlin (1981) showed how the contrast sensitive neurons in the fly's retina maximizes information transmission by adapting to the statistics of its input. How this kind of adaptation of sensoric input can be of use in robotics is presented in section 2.2.

Lungarella et al. (2005) presented a number of information theoretic and other similarity measures to compute relations between sensors. Experiments showed how particular actions of the agent can have an impact on the nature and statistical structure of the robot's sensoric input, where the saliency guided movement decreased the entropy of the input while increasing the statistical dependencies between the sensors.

The present work is also inspired by the more philosophical ideas described by O'Regan and Noe (2001), where the authors propose a new theory that tries to answer the major question of what visual experience is and where it occurs. Many traditional theories rest on the idea that the brain somehow produces an internal image of the world. Instead, O'Regan and Noe (2001) propose that seeing is a way of acting to explore the world mediated by the agent's mastery of the laws of sensorimotor contingencies, which are a set of rules of interdependence between movement and stimulation of the sensors. Some of these ideas have been translated into algorithms by Philipona et al. (2003) and extended in Philipona et al. (2004). There the authors consider the question of whether it is possible for a robot with an unknown body to learn that its body is located in a three dimensional world. They are able to give this question a positive answer by presenting an algorithm that can deduce the dimensionality of the world by analyzing the laws that link motor outputs to sensor inputs. Results from simulation show how a robot with multimodal sensors discovers Euclidean space structure implicit in the equations describing the physical world of the simulation. This is different from the present framework in that the present work does not require continuity for the individual sensors, and, moreover, has also been validated on a real physical robot.

The present article surveys and extends work by Pierce and Kuipers (1997) which was developed further by considering the informational relationships of the sensors in (Olsson et al., 2004b), as will be detailed below. Pierce and Kuipers developed an approach, here called the *sensory reconstruction method*, to create sensoritopic maps from raw and uninterpreted sensors. 'Raw' and 'uninterpreted' here mean that no knowledge of the sensors' or motors' modality, positions, or data representation is given in advance, with all sensor readings normalized between zero and one. A sensoritopic map shows the informational relationships between sensors, where sensors that are informationally related are

close to each other in the maps. In examples, the sensoritopic maps turn out to also reflect the real physical relations and positions of sensors. These maps can then be used by the robot to help it to learn about its sensors and actuators, and to bootstrap learning of motor primitives and control laws. For example, Kuipers et al. (2006) showed how an agent, once relationships between actuators and sensors have been found, can bootstrap learning of homing, path-following, and recognition of places.

Among other steps, Pierce and Kuipers (1997) measured the distances of the temporal signatures of all pairs of sensors. They then projected the sensors into a two-dimensional space in such a way that the projection of sensors match as closely as possible the pairwise distances between their signatures. In their approach a simple one-norm distance measure was used, the average absolute values of difference in sensor values over time.[1] This works well in some cases, but crucially pre-supposes orderedness and continuity in the sensory input. In Olsson et al. (2004b) their work was extended by using the information metric (Crutchfield, 1990), which will be reviewed in section 2.1, to compute the pairwise distances between sensors. Using information theory has several advantages. First of all, it is a well-studied and -understood mathematical framework (Cover and Thomas, 1991). It also removes any assumptions about the ordering or continuity of sensor values, placing sensors together if they are informationally correlated, no matter how far their temporal signatures are in Euclidean space. It also has the advantage of finding both linear and general correlations between sensors, which can be of advantage in, for example, sensor integration (Olsson et al., 2006).

Given a sensory map of the visual field the robot can now explore the relations between its sensors and (so far) unknown actuators by performing motor babbling. While Olsson et al. (2005b) explored this by transforming the continuous two-dimensional map of visual sensors to a discrete array to perform motion flow computations, the assumption of a rectangular sensor field is dropped here. Instead the dimensionality of the sensors is derived and motion flow can be computed for visual fields even if the visual sensors have been discovered to be arranged in non-rectangular layouts, e.g., a circle or a cross. This method is new and can be used to discover and compute information flow in sensors of different modalities. Once motion can be computed, motor babbling can be performed during which the robot develops a repertoire of *sensorimotor laws*, which are rules capturing what effect changes in certain actuator settings will have on its sensory field. Using the sensorimotor laws the robot can then perform basic visually guided movement.

The main aim of the present article is not to design a high performance visual motion tracking system (which the present system is not). It is rather about establishing principles and the use of information theoretic models for autonomous and adaptive sensorimotor development in robotics and to derive a number of phenomena from information theory, for example motion flow.

[1] This is the normalized Manhattan distance between the vectors of sensor values, or, equivalently, the average one-dimensional Euclidean (or Manhattan) distance between sensor readings.

The conceptual power of information-theoretic methods has sometimes been said to be marred by the difficulties they seem to offer to practically relevant implementations, e.g., on robotic systems. Here we present a real physical robotic system that uses information theory to develop sensorimotor control.

The rest of the this article is structured as follows. Section 2 presents the sensory reconstruction method and introduces information theory, entropy maximization of sensor data, creation of sensoritopic maps, and discovery of motion flow in groups of sensors. The next section, section 3, describes the idea of sensorimotor laws and how these can be discovered starting from no knowledge of the structure of sensors or actuators. In section 4 the performed experiments to validate the informational approach to development are presented along with the results. Finally, section 5 concludes the article and presents some ideas for future work.

2 Sensory Reconstruction Method

This section presents the sensory reconstruction method, first introduced in Pierce and Kuipers (1997) and extended by using the information metric (Crutchfield, 1990) in (Olsson et al., 2004b). First a brief introduction is given to information theory in general and the information metric in particular. Then the methods for construction of sensoritopic maps, the grouping of sensors, and the discovery of motion flow in groups of sensors are presented.

2.1 Information Metric

Let $\mathcal{X}$ be the alphabet of values of a discrete random variable (information source) X, where X in the present article will represent a sensor or actuator. The probability that the random variable X assumes a value $x \in \mathcal{X}$ is denoted $P(X = x)$ or, by simplifying the notation, $p(x)$. Then the *entropy*, or uncertainty associated with X is

$$H(X) := -\sum_{x \in \mathcal{X}} p(x) \log_2 p(x) \tag{1}$$

which specifies in bits the information necessary to specify the value assumed by X. The *conditional entropy*,

$$H(Y|X) := \sum_{x \in \mathcal{X}} p(x) H(Y|X = x) \tag{2}$$

$$:= -\sum_{x \in \mathcal{X}} p(x) \sum_{y \in \mathcal{Y}} p(y|x) \log_2 p(y|x) \tag{3}$$

is the uncertainty associated with the discrete random variable Y if the value of X is known. In other words, how much more information one needs to fully determine Y once X is known. The *mutual information* is the information shared between the two random variables X and Y and is defined as

$$I(X;Y) = H(X) - H(X|Y) = H(Y) - H(Y|X). \tag{4}$$

To measure the dissimilarity in the information between two sources the *information metric* (Crutchfield, 1990) can be used. The information metric is the sum of two conditional entropies, or formally

$$d(X,Y) = H(X|Y) + H(Y|X). \tag{5}$$

Note that X and Y in the system are information sources whose $H(Y|X)$ and $H(X|Y)$ are estimated from time series data from each sensor using equation 2 – see also section 2.2.

To better understand the relation between mutual information and the information metric, consider figure 1. Here it is shown that the information metric

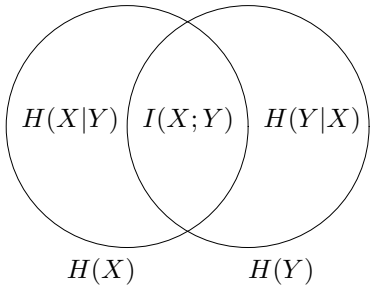

Figure 1: Venn diagram showing the relations between mutual information, $I(X;Y)$, entropy, $H(X)$, and conditional entropy, $H(X|Y)$. The information metric is the sum of $H(X|Y)$ and $H(Y|X)$.

is complementary to mutual information; while mutual information measures what two random variables have in common, the information distance metric measures what they do not have in common.

Many other distance measures exist, e.g, the one-dimensional Euclidean distance used by Pierce and Kuipers (1997), the correlation coefficient, Kullback-Leibler divergence (Cover and Thomas, 1991), Hellinger distance (Basu et al., 1997), and Jensen-Shannon divergence (Lin, 1991), so why the information metric? As discussed in (Olsson et al., 2006), the information metric has a number of interesting and useful properties. First of all is it a proper metric in the mathematical sense.[2] If a space of information sources is metric, it is possible to use some of the tools and terminology of geometry. It is also often useful to

[2]Formally, a metric requires (1) $d(X,Y) = d(Y,X)$, (2) $d(X,Y) = 0$ if and only if $Y = X$, and (3) $d(X,Y) + d(Y,Z) \geq d(X,Z)$. If (2) fails but (1) and (3) hold, then it is a pseudo-metric, from which one canonically obtains a metric by identifying points at distance zero from each other. In the case of the information metric this occurs when X and Y are recoding equivalent, which means that one can map the values of X into Y to determine them completely and *vice versa*. Intuitively speaking, recoding equivalent sources provide the same information. This identification is adopted here and in (Crutchfield, 1990).

be able to talk about sensors in terms of spatial relationships, harnessing spatial intuition. This might be of special importance if the computations are used to actually discover some physical structure or spatial relationships of the sensors, e.g., as in Olsson et al. (2004b), where the spatial layout of visual sensors as well as some physical symmetry of a robot was found by these information theoretic means. This is possible for the visual sensors since physically close visual sensors on average convey similar information and thus have a small information distance. The greater the physical distance is between two visual sensors, the less information they share and hence the information distance is greater. For natural images this shared information between pixels seems to scale as a power-law in the separation distance (Ruderman, 1994).

It is worth noting that the framework of information theory, in this particular case the information metric, provides general methods for quantifying functional relationships between sensors, while many other methods only find linear relationships. For example, a correlation coefficient approaching zero does not imply that two variables actually are independent (Steuer et al., 2002).

2.2 Estimation of Probabilities and Entropy Maximization of Sensor Data

In the work presented here each sensor or actuator variable is as mentioned above modeled as a discrete random variable X that can assume values from a finite alphabet $\mathcal{X}$. The sensor X is a compressed representation of a particular quantity in the world, measured as the value of a sensory variable X_{raw}, that the sensor registers. Thus, a sensor can be viewed as a mapping from a larger alphabet $\mathcal{X}_{raw}$, where $|\mathcal{X}| < |\mathcal{X}_{raw}|$, to the smaller alphabet $\mathcal{X}$, where the number of elements in $\mathcal{X}$ is $N = |\mathcal{X}|$. Perhaps the most common method to do this is by *uniform binning*, where the all possible values of $\mathcal{X}_{raw}$ are divided in to the bins B_i, $\mathcal{X}_{raw} = B_1 \cup \cdots \cup B_N$, where each bin (approximately) contains the same number of values from $\mathcal{X}_{raw}$. That is, $\forall i,\ |B_i| \approx \frac{1}{N}|\mathcal{X}_{raw}|$. Now let each bin B_i correspond to one symbol in the alphabet $\mathcal{X}$ and the probability mass function $p(x)$ can then be estimated from the frequency distribution of the bins.

Even though this method is commonly used is it associated with a number of problems. One example is that the size and placement of the bins might affect whether an important feature of the data is captured or not, and, for example, Lungarella et al. (2005) presents alternative kernel based estimators. In the present article the probabilities are still estimated using a normal frequency count method, but instead of uniform binning *adaptive binning* is used, also known as *entropy maximization* or *sampling equalization*. Adaptive binning is most easily explained using an example.

Consider the grey-scale image depicting Claude Shannon in figure 2(a) taken from (Olsson et al., 2005c), where each pixel can have a value between 0 and 255, meaning that $\mathcal{X}_{raw} = \{0, 1, \ldots, 255\}$. How can this image be compressed to only 5 different pixel values, $\mathcal{X} = \{0, 1, \ldots, 4\}$? Figure 2(b) shows the image compressed using uniform binning and, as seen in figure 2(e), the distribu-

(a) No binning (b) 5 uniform bins (c) 5 adaptive bins

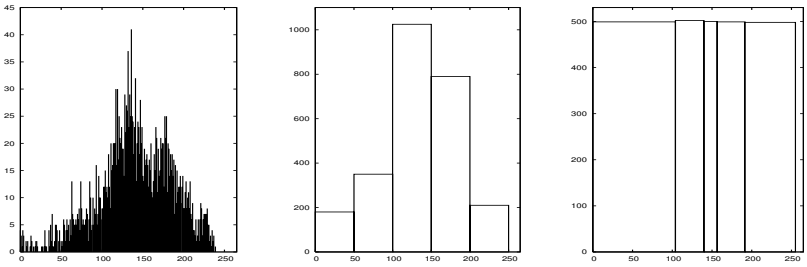

(d) No binning - histogram (e) 5 uniform bins - histogram (f) 5 adaptive bins - histogram

Figure 2: Figure 2(a) shows a 50 by 50 pixels grey-scale image of the founder of information theory, Claude Shannon, and figure 2(d) the corresponding histogram of pixels between 0 and 255. Figure 2(b) shows the same picture where the pixel data (0-255) is binned into only 5 different pixel values using uniform binning and figure 2(e) the frequency of pixel values. Finally, figure 2(c) shows the same picture with the pixel data binned into 5 bins using adaptive binning and figure 2(f) the corresponding histogram [after (Olsson et al., 2005c)].

tion of grey-scales in figure 2(a) is not uniform, with most pixels in the range $\{100, 101, \ldots, 200\}$. The entropy of the encoded image is approximately 1.97 bits, which is less than the maximal theoretical entropy of $\log_2 5 \approx 2.32$ bits. From an information theoretical point of view this means that this encoding is non-optimal since the entropy of the encoded image is less than the maximal possible entropy of the image. Now, consider figure 2(c) which also uses 5 bins, where (at least if studied from a distance) the image subjectively seems to convey more detail about the original image. Here the original values have been binned in such a way that each bin contains approximately the same number of pixels, which means that the entropy of the image is close to the maximum of $\log_2 5 \approx 2.32$ bits. This can also be considered as adding more resolution where

most of the data is located.

Similarly, in nature it has been found that the contrast encoding in the visual system of the fly is adapted to the specific statistics of it environment (Laughlin, 1981), such that, just as in the image of Claude Shannon above, the entropy of the graded response is maximized. More formally, given a raw sensor X_{raw}, the aim is to find a partitioning of the data into the N bins of the alphabet $\mathcal{X}_{raw} = B_1 \cup \cdots \cup B_N$ such that each bin B_i is equally likely given a certain set of data. That is,

$$P(X_{raw} \in B_i) \approx \frac{1}{N}, \tag{6}$$

which is equivalent to the entropy of which bin data falls into being maximized. Another way to view this is to make the distribution as close as possible to equidistribution, given a particular set of data. It is important to note that while entropy maximization changes the entropy of individual sensors, it still maintains correlation between sensors. This was shown in Olsson et al. (2005c) where adaptive binning was applied to integration of visual sensors of different modalities.

The experiments performed by Brenner et al. (Brenner et al., 2000) also indicate that this kind of entropy maximization constantly is happening in the motion sensitive neurons of the fly. This can be implemented by estimating the probability distribution of the most recent T time steps and changing the transfer function accordingly. In the experiments performed in the present article this algorithm is implemented using frequency distribution estimators to estimate the probability distributions. Since multiple instances of the same value of X_{raw} are added to the same bin, the distribution need not be completely uniform, as in figure 2(f). The sliding window in the present implementation does not use decay, which means that more recent values do not affect the distribution more than older ones within the window.

2.3 Creating Sensoritopic Maps

In the sensory reconstruction method (Pierce and Kuipers, 1997; Olsson et al., 2004b), sensoritopic maps are created that show the informational relationships between sensorimotor variables, where sensors that are informationally related are close to each other in the maps. In an appropriate environment, the sensoritopic maps are likely to also reflect the real physical relations and positions of sensors. For example, if each pixel of a camera is considered a sensor, is it possible to reconstruct the organization of these sensors even though nothing about their positions is known. Sensoritopic maps are often depicted as two-dimensional for visualization reasons, but it is possible to find more appropriate dimensions of the sensors, see section 2.4, and also to create maps of higher dimensions. Figure 3 shows an example of a sensoritopic map for a SONY AIBO robot.

To create a sensoritopic map the information distance is first computed between all pairs of sensors as described above in section 2.1. The resulting distance matrix specifies the informational distance between all pairs of sensors.

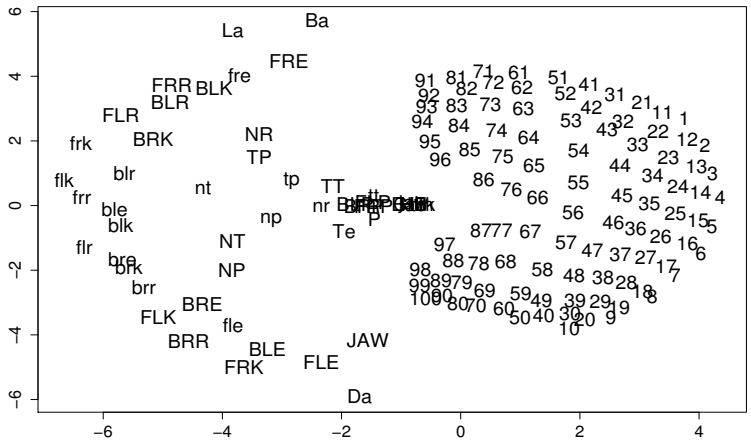

Figure 3: A sensoritopic map created by the sensory reconstruction method using the information metric. In this example there are 150 sensors, including 100 image sensors that are labeled 1-100 to the right in the map. The distance unit is bits [after (Olsson et al., 2004b)].

To create a sensoritopic map from the distance matrix the N sensors are projected to a plane. In general, a perfect mapping given a N by N distance matrix requires $N-1$ dimensions, but often a lower dimensionality is the most appropriate to reflect the nature of the sensors and the data. For example, a rectangular visual field is best represented in two dimensions, and that this is the case can actually be shown algorithmically using only the sensor data – see section 2.4 and figure 6 in the experimental section.

To project the sensors of the distance matrix to a lower dimension is a constraint-solving problem that can be solved by a number of different methods such as metric-scaling (Krzanowski, 1988), Sammon mapping and elastic nets (Goodhill et al., 1995), as well as self-organizing maps (Kohonen, 2001; Polani, 2002). Here it was decided after some experimentation to use the relaxation method presented in (Pierce and Kuipers, 1997), primarily because of its speed and qualitatively better results found by visual inspection of reconstruction of known visual layouts.

In the relaxation method each sensor $S_i, 1 \leq i \leq N$, is assigned a random position $\mathbf{p_i}$ on a plane. The force $\mathbf{f_i}$ on each point $\mathbf{p_i}$ is

$$\mathbf{f_i} = \sum_{j=1}^{N} \frac{(||\mathbf{p_i} - \mathbf{p_j}|| - d(X_i, X_j))(\mathbf{p_j} - \mathbf{p_i})}{||\mathbf{p_j} - \mathbf{p_i}||}. \tag{7}$$

Each point $\mathbf{p_i}$ is moved according to the force $\mathbf{f_i}$ acting on that point:

$$\mathbf{p_i} = \mathbf{p_i} + \frac{\mathbf{f_i}}{N} \tag{8}$$

The algorithm is iterated for a number of time steps moving each point according to equation 8 until only very small or no movements occur, which usually takes between 2000 to 3000 iterations for 100 sensors.

2.4 Grouping of Sensors and Dimensionality

Now sensors that are close can be grouped together. Two sensors, x_i and x_j, are similar according to Pierce and Kuipers (1997), if

$$x_i \approx x_j \quad \text{iff} \quad d(x_i, x_j) \leq \min\{\epsilon_i, \epsilon_j\}, \tag{9}$$

where the ϵ_i is calculated from the minimum distance, $\epsilon_i = 2\min_j\{d(x_i, x_j)\}$, to any of neighbours. To form equivalence classes, (Pierce and Kuipers, 1997) use the *related-to* relation which is the transitive closure of the similarity relation $\approx$ This is computed recursively via

$$i \sim j \quad \text{iff} \quad i \approx j \vee \exists k : (i \sim k) \wedge (k \sim j). \tag{10}$$

Given such an equivalence class its dimension can be computed. Let $\sigma^2(m)$ be the variance accounted for by dimension m. Then the right number of dimensions can be found by maximizing $\sigma^2(m) - \sigma^2(m+1)$. For example, for a set of visual sensors arranged in a two-dimensional grid $m = 2$. To compute $\sigma^2(m)$, metric-scaling (Krzanowski, 1988) can be used, where the eigenvalues produced in the calculation for each dimension provide a measure of the variance – see section 4.1 for an example.

2.5 Discovering Motion Flow

It is often the case that a robot is in an environment where movement has a piecewise smooth effect on the sensory data produced by some of the sensors. For example, most real world scenarios have this effect on visual sensors unless there is complete darkness or the visual input is random. The same is also true for, e.g., infra-red sensors that measure the distance to the closest obstacle, and in some cases, for tactile sensors. This effect can be exploited by a robot to perform sensory-guided movement by computing motion flow in the sensors which can tell the robot about the velocity and direction of movement, caused by the robot itself or movement of objects in the environment.

Traditional motion flow techniques, e.g., optical flow analysis (Lucas and Kanade, 1981), are based on the idea that for most points in the image, neighbouring points have approximately the same brightness. In other words, the world is made up of continuous objects over which brightness varies smoothly. In Olsson et al. (2005a), this condition was relaxed and generalized by assuming that neighbouring sensors are more highly informationally correlated, which

does not necessarily entail that they have similar values. This can, e.g., be the case in a visual field where some sensors represent the red channel, while other sensors in other positions represent the green or the blue channel. In (Olsson et al., 2005a) the motion flow in a certain sensor S_t^x was computed by time-shifting the sensors around that sensor comparing the informational distance between S_t^x and other sensors, $S_{t-\Delta t}^y$, in different positions and times. Given that a sensor with a certain time-shift is found to be closely or completely correlated with the sensor for which the motion flow is to be computed, the speed and direction of the flow can be computed from the sensoritopic map.

This method has several drawbacks. One is under-sampling. While ordinary optical flow computation can be done using only two frames of data, much more data is needed for this information theoretical method. It is also computationally expensive to compute the entropies for a large number of sensors with different time-shifts. One advantage, though, is the fact this method does not assume a discrete and rectangular visual field, something that is often taken for granted in normal optical flow algorithms. It should also be noted that this method does not attempt to outperform other methods. Instead, as mentioned before, the aim is to guide development and derive generalized phenomena from information theory.

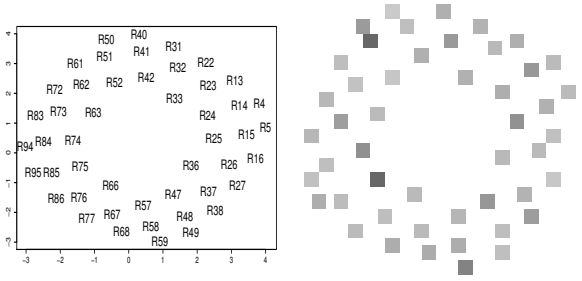

(a) Circular Visual Field

(b) Image for optical flow calculation

Figure 4: Figure 4(a) shows the reconstructed visual field and 4(b) the corresponding image used in the optical flow calculation.

In the present article sensor fields can be of any shape, not just rectangular and two-dimensional grids, which makes this last property of the information theoretic motion flow method attractive. Thus, it was decided to use traditional optical flow calculations but combined with the idea that the visual field not have to be of a rectangular shape. This is implemented as follows. Consider that the sensory reconstruction method has found that the visual field of a robot is two-dimensional and that the sensors are placed in a circle in the sensoritopic map, see figure 4(a). For each frame of data an image is created which is initially

of uniform colour, e.g., all white. By discretizing the continuous coordinates of the sensors in the sensoritopic map of 4(a), the value of each particular sensor can placed in its corresponding position in the image, resulting in an image such as figure 4(b). Now the optical flow can be computed between two images[3] and the motion vectors in the pixels corresponding to sensors can be saved.

How can the optical flow technique be applied to sensor fields with sensors of different modalities, e.g., the red, green, and blue channels of camera pixels? Instead of the actual values varying smoothly over the sensors, just the correlatedness between the sensors will vary smoothly, and the closer two sensors are in the visual field, the smaller the informational distance will be. This may result in a lack of smooth change in sensory values as one moves over the sensoritopic map of the visual field. How to handle this effectively is an open problem with a number of possible candidate solutions. One possible solution may be to find a relabeling of the sensor data where the sensor values vary smoothly, while keeping the same information distance between all pairs of sensors as before. Another possibility is to partition the sensors based on value similarity and calculate the optical flow separately for each partition.

3 From Motor-babbling to Structured Movement

In this section it is described how learning via motor babbling can be performed once the structure of the sensors has developed. First sampling of actuator space is discussed in section 3.1, before how to learn sensorimotor laws in section 3.2, and how to use the learned laws to perform actions in section 3.3.

3.1 Sampling Actuator Space

The goal of motor babbling is to sample the space of possible actuator settings in an as exhaustive way as possible and learn how to associate actions with sensory perceptions. While the methods presented in section 2 of the sensory reconstruction method were general and not dependent on a particular embodiment, sampling of actuator space is more dependent on the hardware used. For example, some actuator settings might physically harm the robot, or cause legs to get entangled in each other. Thus, some constraints might be necessary when sampling the possible settings of a robot's actuators.

In general, the aim is to sample, in as much detail as possible, the set of possible actuator settings, to understand how all the different actuator settings affect the sensors. As a simple example, consider a robot with two actuators, a_1 and a_2, where $a_1, a_2 \in \{-1, 0, 1\}$. This means that the whole actuator space consists of only 9 possible settings. In most robots this space is much larger, with more actuators, and more possible values (maybe continuous) for each actuator. There might also be a many-to-one mapping between the actuator

[3]This is implemented by the authors using the Lucas-Kanade algorithm (Lucas and Kanade, 1981) implemented in OpenCV:
http://www.intel.com/research/mrl/research/opencv/

vectors and the actual values sent to the motor, which means that more than one actuator vector might give the same effect in the motors.

Also, for more robust learning each movement needs to be performed several times, resulting in a number of samples of sensor data for each possible movement. This is necessary since many sensors are sensitive to noise and the structure of the environment. There might also be moving objects in front of the robot, something that will seriously affect for example motion flow calculations.

3.2 Learning Sensorimotor Laws

The result from motor babbling as described above is a collection of sensor data associated with each possible actuator setting. This data can be used to compute the *sensorimotor laws*. Here a sensorimotor law is defined as the average effect a certain actuator setting will have on a set of sensors.

Let the sensor S_n at position (x, y) in the two-dimensional visual field of section 2.5 be associated with its average effect for a certain actuator setting a as $\langle \vec{v}_{x,y,a} \rangle$. The average effect is defined as the average optical flow in this sensor calculated as described in section 2.5 using the Lucas-Kanade optical flow algorithm (Lucas and Kanade, 1981) over all frames of data T for that particular setting a, is

$$\langle \vec{v}_{x,y,a} \rangle = \frac{\sum_{t=1}^{T} \vec{v}_{x,y,a,t}}{T}, \tag{11}$$

where $\vec{v}_{x,y,a,t}$ is the optical flow in $\vec{v}_{x,y,a}$ between time t and $t+1$.

A sensorimotor law for actuator setting a, SM_a, is defined as a matrix consisting of the average motion flow vectors for all the sensors in the two-dimensional grid with width w and height h:

$$SM_a = \begin{pmatrix} \langle \vec{v}_{1,1,a} \rangle & \langle \vec{v}_{2,1,a} \rangle & \cdots & \langle \vec{v}_{w,1,a} \rangle \\ \langle \vec{v}_{1,2,a} \rangle & \langle \vec{v}_{2,2,a} \rangle & \cdots & \langle \vec{v}_{w,2,a} \rangle \\ \vdots & \vdots & \ddots & \vdots \\ \langle \vec{v}_{1,h,a} \rangle & \langle \vec{v}_{2,h,a} \rangle & \cdots & \langle \vec{v}_{w,h,a} \rangle \end{pmatrix} \tag{12}$$

Since the actual field of the robot might not be rectangular, see section 2.5, some motion vectors will be of zero length since they are not associated with a sensor. This is not a problem since the layout of the visual field is the same for all sensorimotor laws SM. It might be the case that different actuator settings will have the same, or very close to the same, effect on the sensors of the robot, even though they are different for an external observer. These two actuator settings should still be combined into one sensorimotor law as the only information the robot has about itself and its environment is its own sensors.

For many movements if a rectangular visual field is used, e.g., pan and tilt motions, the vectors of the vector field will be of more or less uniform magnitude. If all possible actuator settings are of this kind in a certain experiment, sensorimotor laws can be simplified by computing the average over the whole matrix,

excluding positions not corresponding to any sensor, and let the sensorimotor SM_a be equal to the average of all motion vectors $\vec{v}_{x,y,a}$.

3.3 Sensory-Guided Movement

The set of sensorimotor laws can now be used to perform actions where a specific effect on the sensors is desired. For example, if the robot perceives a motion the motion can be tracked by performing any action from the set of sensorimotor laws that has as similar as possible an effect on the sensors as the observed motion.

To track a particular motion, the robot needs to find the sensorimotor law with effect closest to the perceived movement P, represented as an optical flow matrix like the sensorimotor law, equation 12. Given a set of learned sensorimotor laws, a law closest to the perceived visual movement can be found by minimizing some measure $D(SM, P)$ between the perceived flow P and sensorimotor law SM. Perhaps the most natural is the one-norm:

$$D(SM, P) = \sum_{x,y} |SM_{x,y} - P_{x,y}| \tag{13}$$

If the perceived motion is caused by a moving object much smaller than the visual field, many of the vectors in P will be zero. Thus, vectors that are of zero magnitude in P are not compared. An alternative to the one-norm is presented by Lee and Seung (2001), which is computed separately for the x and y component of the flow matrix. This distance is defined as

$$D(SM^c||P^c) = \sum_{x,y} \left(SM^c_{x,y} \log_2 \frac{SM^c_{x,y}}{P^c_{x,y}} - SM^c_{x,y} + P^c_{x,y} \right), \tag{14}$$

where c is either the x or y component of the flow. This distance measure reduces to the Kullback-Leibler divergence if $\sum_{x,y} SM_{x,y} = \sum_{x,y} P_{x,y} = 1$ and is not symmetric, but it can be made symmetric by $D(SM^c, P^c) = D(SM^c||P^c) + D(P^c||SM^c)$.[4] Finally to compute the measured difference between the two matrices P and SM the distances of the two components x and y are added together. The methods of this section generalize in the obviously manner from sensory fields of two-dimensions to n-dimensional ones.

4 Experiments

In the experiments presented below to validate the framework, a SONY AIBO[5] robotic dog was used. The experiments were performed in a laboratory with no windows that could seriously affect light conditions. The robot was placed in a sitting position overlooking two desks, see figure 5.

[4] It should be noted that strictly speaking, this is not a metric since the triangle inequality does not hold for the Kullback-Leibler divergence (Cover and Thomas, 1991).

[5] AIBO is a registered trademark of SONY Corporation.

Figure 5: The experimental setup with the robot sitting in the lab.

The AIBO has many degrees of freedom and to simplify the experiments we decided to only use the pan (horizontal movement) and tilt (vertical) motors of the head of the AIBO. To put less strain on the neck of the AIBO (which is the component that most often breaks down), we let the maximum pan and tilt be 0.5 and the minimum -0.5, where the possible values of the neck are in the range $[-0.6, 0.6]$. To move the neck the desired position of the pan and tilt motor were set 100 times per second, adding the value of the corresponding actuator setting each time step, until the maximum or minimum value of the pan or tilt had been reached. In our experiments the possible desired speed, s, for each of the two actuators was limited to three different values, $s \in \{-0.02, 0.0, 0.02\}$.

The 88 by 72 pixel visual image from the AIBO was downsampled to a 10 by 10 image by averaging 20 times per second. Each pixel had a resolution of eight bits per channel (red, green, blue), in the range $\{0, 1, \ldots, 255\}$. For the initial experiments described in section 4.1 and 4.2, 10 runs each collecting 6000 frames of data were performed. In these runs the robot moved its head in one of the eight possible directions until the limit for either the tilt or pan was reached, then another one of the possible movements from that position was chosen at random. To compute the information metric of section 2.1 10 bins were used and the window size of the adaptive binning – see section 2.2, was 100.

4.1 Reconstruction of Visual Fields

In previous work (Olsson et al., 2004b, 2005b,a) it has always been assumed that the visual field is a square array. Here results are presented for different subsets of the whole visual field, showing the reconstruction of visual fields of the different shapes shown in figure 7(a), 7(b), and 7(c). Also a one-dimensional row of vision sensors was used as a sensor field. The input to the robot was the sensor data from the red sensors of the visual field without any knowledge of the sensors' positions. All the 10 sets of 6000 frames of data were used in each case and they all produced similar results, and the results presented here are

typical examples.

The first problem is what dimensionality the sensor field should be represented in. This problem is solved using the method of section 2.4, where the eigenvalues from metric-scaling are used to find the cut-off dimension which contributes the most variation in the sensor data.

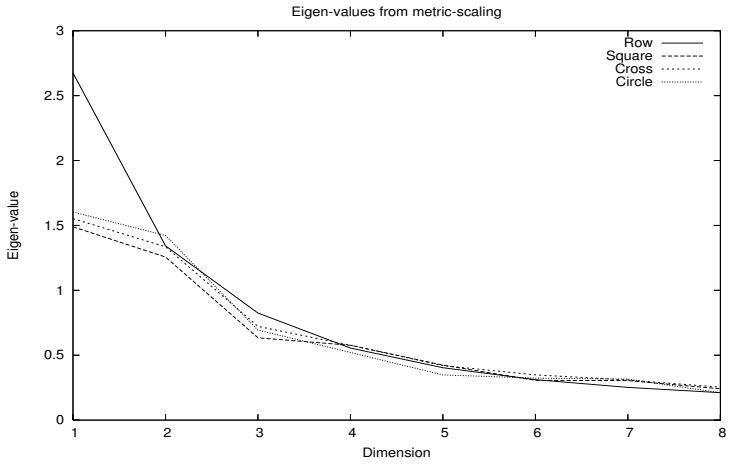

Figure 6: Figure shows the eigenvalues, which correspond to the variance, per dimension from metric-scaling. The eigenvalues are normalized by dividing by the number of sensors for each layout. For the one-dimensional row of sensors, the biggest difference in variance is between dimension 1 and 2, which means that these sensors are best represented in one dimension – see section 2.4. In the variance for the cross, circular, and square layouts, the biggest difference is between the second and third dimension, meaning that two dimensions is the most appropriate.

Figure 6 show the eigenvalues for the AIBO data from the three layouts in 7(a), 7(b), 7(c), and also for a visual field consisting of one row of 10 visual sensors. The figure shows how the biggest variance difference for the one-dimensional row of sensors is between dimension one and two, meaning that the data correctly should be one-dimensional. For the cross, circle, and square, on the other hand, the biggest variance is between dimension two and three, correctly indicating that this data is from a two-dimensional visual field.

Given the dimensionality, the sensory reconstruction method can now be used to reconstruct the visual field. Figure 7 shows the resulting reconstructed visual fields for the cross, circular, and square layouts. It is important to note that the method used to project the sensors in two dimensions, the relaxation algorithm, initially places each sensor in a random position. Thus, each time a map is created it will look different. But, since the method does not aim at recreating the actual physical positions, but rather the spatial and informational relations between the sensors, this is not an issue. It is also the case that the

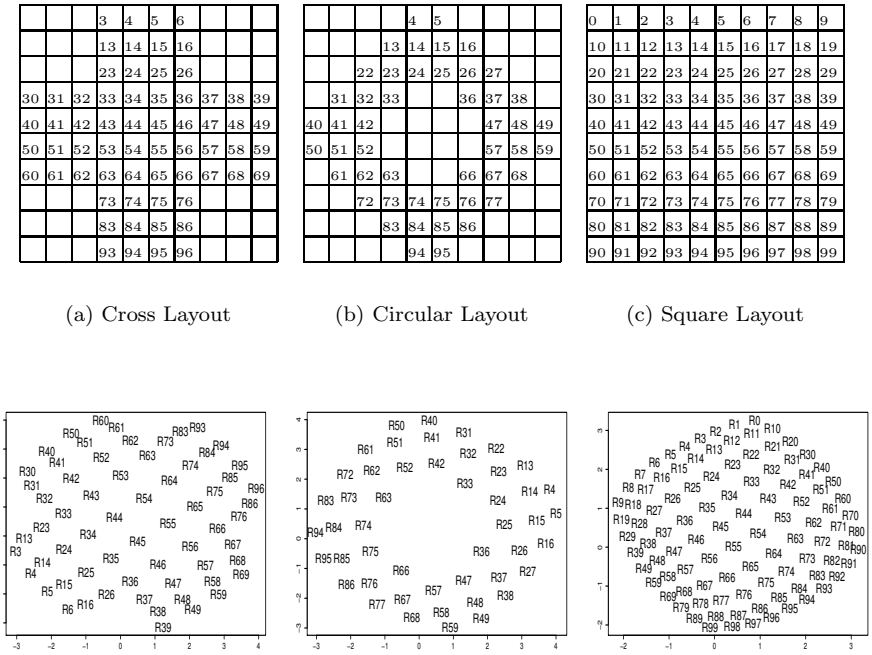

(a) Cross Layout (b) Circular Layout (c) Square Layout

(d) Cross Visual Field (e) Circular Visual Field (f) Square Visual Field

Figure 7: Layout of the sensors, figure 7(a) to 7(c), and the reconstructured visual fields represented as sensoritopic maps, figure 7(d) to 7(f).

real orientation of the field can not be found without higher level visual processing. Consider, for example, figure 7(d), where the field is roughly 245 degrees rotated and also the mirror image of the layout of figure 7(a). None of this matters from the developing agent's perspective, as it can use the constructed sensoritopic map for motor learning regardless of an external observer's notion of "real orientation".

One can also look at how the sensoritopic maps develop over time as more statistical structure is found. Figure 8 shows how the square visual field developed from 10 to 1000 frames of data for a particular run. After only 10 frames of data there is not much structure. With more frames of data the sensors located closer to each other in the visual field decrease their informational distance, and the sensoritopic map more closely resemble the real visual field. The examples shown in this paper use a small visual field of 10 by 10 sensors for visual clarity in the figures but the authors have experimented with visual fields up to fifty by fifty sensors without qualitatively different results. It should also be noted that this sensoritopic reconstruction can be done in real time.

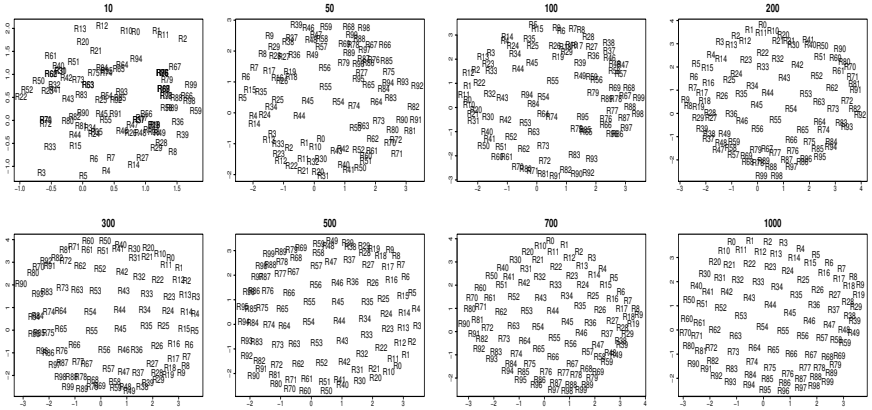

Figure 8: Development of sensoritopic map of visual field over time from 10 to 1000 time steps.

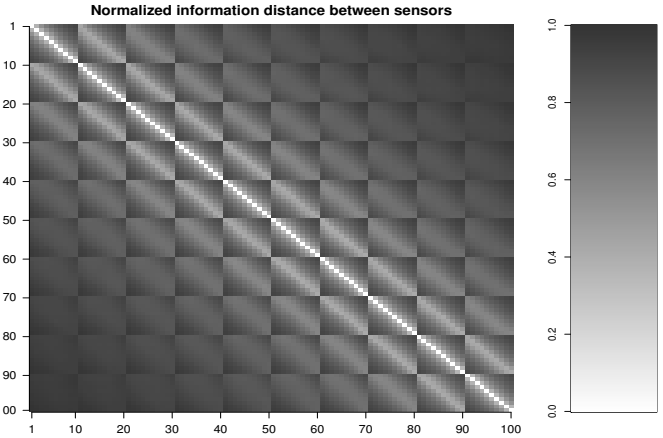

Figure 9: Normalized information distances between all pairs of sensors for the square visual field of 7(c) with 0 and 1 visualized as white and black, respectively. The distance between a sensor and itself is always zero, hence the white diagonal line. The further away two sensors are the less are they correlated and the greater is the information distance.

How is it that the relative positions of physical sensors can be found just by computing the informational distances between the sensors? The key is that the closer two sensors are in the physical visual field, the smaller the informational distance is between the two sensors. Figure 9 shows the distance matrix for all sensors from the square visual field of figure 7(c). The distance between a

sensor and itself is zero and the 10 by 10 squares in the matrix indicate that the visual field is 10 by 10 sensors. Looking closer at the data one finds that the information distance between the sensors scales as a power-law with regards to the physical distances between the sensors. This is in line with the statistics of natural images found by, for example, Ruderman (1994), where it was shown how the mutual information between pixels scales as a power-law with regards to the physical distance between the pixels in images from natural environments.

Up until now the term "reconstructed" has been used in an informal way, where a visual field is reconstructed if the sensoritopic map and the real layout of the sensors look similar, see, e.g., figure 7. One way this similarity can be formally quantified is by computing the relative distances between pairs of sensors in the reconstructed visual field and the real layout of the sensors. Let $r_{i,j}$ be the Euclidean distance between two sensors i and j in the reconstructed map, and $\ell_{i,j}$ the distance between the same two sensors in the real layout, where the x and y coordinates in both cases have been normalized into the range $[0.0, 1.0]$. Now the average distance between all pairs of sensors can be compared,

$$d(r, \ell) = \frac{1}{N^2} \sum_{i,j} |r_{i,j} - \ell_{i,j}|, \tag{15}$$

where N is the number of sensors. This compares the relative positions of the sensors and not the physical positions, and $d(r, \ell)$ will have a value in the range $[0.0, 1.0]$. A distance of zero means that the relative positions are exactly the same, and sensors placed at completely random positions will have an average distance of approximately 0.52.

For each of the three layouts in figure 7, the distance was computed between the real layout and 1000 reconstructed sensoritopic maps created for each layout. For the square layout the average was 0.051 (st.d. 0.012) and for the cross layout 0.075 (st.d. 0.016), and for the circle layout 0.057 (st.d. 0.013). This should be compared with randomly generated maps, where the average distance is significantly different at approximately 0.52 (st.d 0.002). This measure cannot tell in absolute terms whether a sensoritopic map really is a reconstruction of the real visual layout, but values of the average distance close to zero definitely means that much of the structure in the real layout has been captured.

4.2 Reconstruction of Visual Fields of Different Modalities

Now imagine a visual field consisting of sensors of different modalities such as in figure 10(a), where some sensors report the red channel, others green, or the blue channel of a pixel. This layout is unknown to the robot and just as before the input is 100 streams of the data and the problem is to reconstruct the visual field, even though the sensors are of different visual modalities.

Using the same visual AIBO data as before, sensoritopic maps were created for all 10 sets of data. Reconstruction was successful for all 10 sets, with an

R0	R1	R2	R3	R4	R5	R6	R7	R8	R9
R10	R11	R12	R13	R14	R15	R16	R17	R18	R19
R20	R21	R22	R23	R24	R25	R26	R27	R28	R29
R30	R31	R32	G33	G34	G35	G36	G37	G38	G39
G40	G41	G42	G43	G44	G45	G46	G47	G48	G49
G50	G51	G52	G53	G54	G55	G56	G57	G58	G59
G60	G61	G62	G63	G64	G65	B66	B67	B68	B69
B70	B71	B72	B73	B74	B75	B76	B77	B78	B79
B80	B81	B82	B83	B84	B85	B86	B87	B88	B89
B90	B91	B92	B93	B94	B95	B96	B97	B98	B99

(a) Separated layout with (R)ed (G)reen and (B)lue sensors in different positions

S0	S1	S2	S3	S4	S5	S6	S7	S8	S9
S10	S11	S12	S13	S14	S15	S16	S17	S18	S19
S20	S21	S22	S23	S24	S25	S26	S27	S28	S29
S30	S31	S32	S33	S34	S35	S36	S37	S38	S39
S40	S41	S42	S43	S44	S45	S46	S47	S48	S49
S50	S51	S52	S53	S54	S55	S56	S57	S58	S59
S60	S61	S62	S63	S64	S65	S66	S67	S68	S69
S70	S71	S72	S73	S74	S75	S76	S77	S78	S79
S80	S81	S82	S83	S84	S85	S86	S87	S88	S89
S90	S91	S92	S93	S94	S95	S96	S97	S98	S99

(b) Integrated layout with three sensors, $S = \{R, G, B\}$, in each position

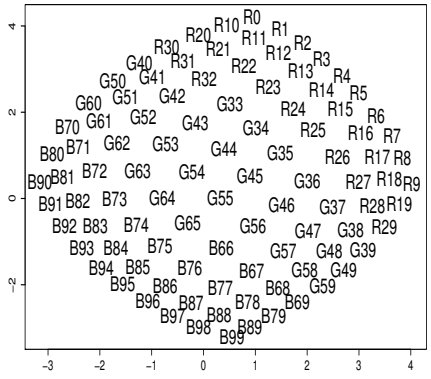

(c) Red, Green, Blue separated

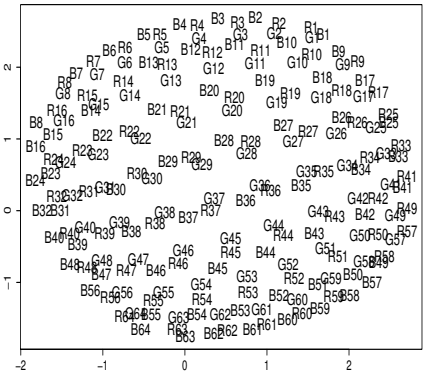

(d) Red, Green, Blue integrated

Figure 10: Sensory reconstruction with sensors of different modalities. Figure 10(a) shows the layout of the red, green, and blue sensors in a separated visual field and figure 10(c) the reconstructed visual field. Figure 10(b) shows the layout of an integrated visual field with a total of 300 sensors with one red, one green, and one blue sensor in each position. Figure 10(d) shows the integrated reconstructed visual field.

average distance between the real layout and the sensoritopic maps – see equation 15 in section 4.1 – of 0.061 (st.d. 0.008). Figure 10(c) shows an example

of the reconstructed visual field. Since the red, blue, and green sensors may have different ranges of data, the information metric together with the entropy maximization technique make this possible – see section 2.2.

A related problem is to consider each pixel as having three sensors, one for the red channel, one for the green channel, and one for the blue channel, giving a total of 300 sensors. This was first studied in (Olsson et al., 2005c), where the problem was to find what sensors that correspond to the same location in the image, even though they are of different modalities. Figure 10(d) shows the reconstructed map from the same AIBO data as used above applied to this problem. All 10 data sets resulted in similar reconstructions. This result is due to the fact that the information metric finds nonlinear correlations, see (Olsson et al., 2006). In this case the average distance (equation 15) between all 300 sensors for all 1000 sensoritopic maps created for each of the 10 data sets was 0.092 (st.d. 0.019).

This is an example of autonomous sensory fusion where sensors of different modalities are combined. A well-studied example of this in neuroscience is the optic tectum of the rattlesnake, where nerves from heat-sensitive organs are combined with nerves from the eyes (Newman and Hartline, 1981).

4.3 From Unstructured Sensors and Actuators to Visually Guided Movement

The aim of this next experiment was to have the robot develop from no knowledge of its sensors and actuators to basic visually guided movement. As described above, 100 visual sensors were used and the robot had two actuators, pan and tilt.

Once the visual field has been reconstructed as shown in section 4.1 is it possible to compute the optical flow in the sensors as described in section 2.5. Given that the robot has developed a sensoritopic map of its visual field so that it can compute optical flow, it is possible to start to experiment with how different settings of the actuators affect the sensors. As described in section 3.1, this body babbling is performed by sampling the space of possible actuator settings. As mentioned above this particular experiment only used three possible settings per actuator, giving a total of eight possible non-zero settings. Each possible setting was executed 50 times, giving a total of roughly 500 frames of data per setting.

Given this data the sensorimotor laws were computed as per section 3.2. Figure 11 shows the computed sensorimotor laws. As can be seen the directions of the motions' vectors do not correspond exactly to an external observer's notion of direction of the movement. This is due to the reconstructed field being angled with respect to the x and y axes, which is unavoidable since the robot cannot know anything about the real locations of the sensors, and grounds all actions and knowledge about its sensors and actuators in its own sensorimotor perceptions.

To verify that the method computed similar sensorimotor laws each time the experiment was performed 10 times with each movement being executed

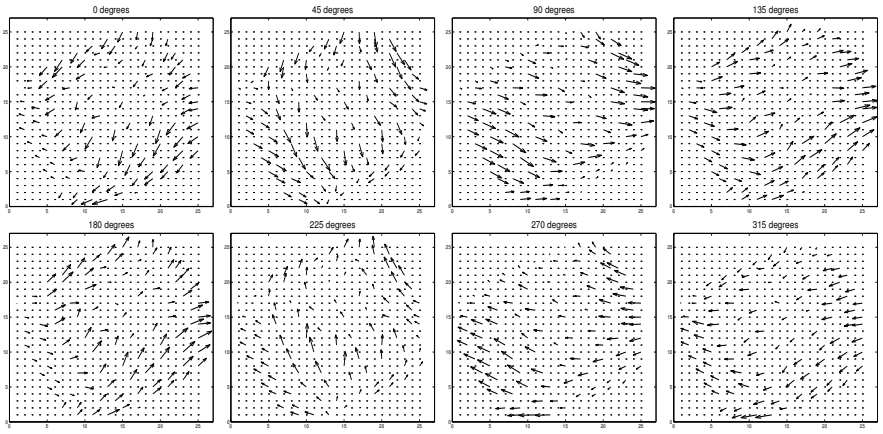

Figure 11: The discovered sensorimotor laws corresponding to motion of the head where 0 degrees is moving the head up, 180 degrees moving the head down, and 90 degrees to the right.

50 times as described above. Then the two distances functions of section 3.3 were used to compute the distance between the sensorimotor law computed for each time. The results showed that the learned laws were very similar, with an average distance of 1.2% for the one-norm distance and 0.9% for equation 14 averaged over all eight actuator settings and the 10 experiments.

When the sensorimotor laws have been learned, the robot can now use these laws to perform movement guided by its sensors. To verify this step two experiments were performed. In both experiments the AIBO was placed in front of a wall of uniform colour with a video projector displaying images on the wall. In the experiments the whole image on the wall moved in different directions. This was implemented by having an image, shown in figure 12, larger than the screen and only displaying a part of the image at any given time. The speed of the moving image was adjusted to be similar to the speed of the moving AIBO head. When the image moved the robot used two or 10 frames of data to compute the optical flow to find the sensorimotor law most similar to the experienced movement. In the case of ten frames the optical flow was averaged as in equation 3.2. The motion closest to the perceived motion could then be compared with the actual movement of the image on the wall. Table 1 shows a comparison of the different distance measures and the performance of the motion detection for 200 examples of each movement. As expected the averaged motion flow over 10 frames was more often correctly classified. The one-norm and Kullback-Leibler distance measures performed similarly, with a slightly higher average performance for the latter, equation 14, for the optical flow computation between only two frames.

Figure 12: Image displayed on wall in front of robot.

Action \ $D(SM, P)$	O. 2	O. 10	K. 2	K. 10
0 degrees	75.5%	90.1%	79.2%	91.9%
45 degrees	77.0%	92.8%	78.7%	92.4%
90 degrees	76.6%	91.2%	79.7%	94.4%
135 degrees	75.1%	93.0%	77.7%	93.2%
180 degrees	74.8%	91.5%	75.3%	92.6%
225 degrees	77.1%	93.3%	77.2%	93.8%
270 degrees	76.1%	93.2%	78.8%	92.5%
315 degrees	74.8%	91.4%	76.6%	93.0%

Table 1: Percentage of correctly detected motion flows using the (o)ne-norm and (K)ullback-Leibler measure for 2 and 10 frames for the different directions out of eight possible flows.

5 Conclusions

The present article has presented methods that allow a robot to develop from no knowledge about its sensors and actuators to performing basic visually guided movement. Central to the method is the information metric (Crutchfield, 1990), which measures the informational distance between two sensors. This enables the robot to find correlations in sensors which are not linear. Information theory is often viewed as having a strong conceptual power but also having numerical and statistical requirements that makes it problematic to use information theory in practical implementations of, for example, robotic control systems. The methods presented here can all be used in online learning systems executing in real-time.

To develop from unknown sensors and actuators the robot starts by performing random movements. As the robot performs the movements the informational

relations of the sensors unfold (see figure 8). Given a set of sensors their dimensionality can be computed and a sensoritopic map can be created using the sensory reconstruction method. The sensoritopic map reflects the informational geometry of the sensors, where sensors that are highly correlated are close in the map. Such proximity need not respect sensory modality, as occurs, e.g., in the optic tectum of rattlesnakes, where bimodal neurons receive input from the retina as well as infrared-sensing pit organs, and the neurons cross-modally integrate sensory information (Newman and Hartline, 1981). Given this map a method was presented for computing optical flow in reconstructed visual fields that are not rectangular. It was also shown how this method can be extended to compute motion flow in a visual field consisting of sensors of different modalities, for example a mix of red, green, and blue sensors. When optical flow can be computed the average effect of different actuators settings can be learned. The method of sensory field reconstruction and sensorimotor learning could also be applied to sensory fields other than two dimensional visual ones – for example tactile sensing over the skin or pressure changes in volume in which sensors are distributed. This enables the robot by motor babbling to build a library of sensorimotor laws which specify how its possible actions affect its sensors. Using these laws the robot can see motion flows in its visual field and then perform a movement which will have a similar effect. This can be used for basic imitation or motion tracking.

The underlying principle of the presented work is to guide development by informational means. Apart from the advantage of finding nonlinear relations between sensors, information theory also has the advantage of being modality and dimension independent. This is something that will be explored in future work where the same methods will be applied to infrared sensors configured in a three dimensional space. The presented method can also be used to investigate what a robot in principle can know about its environment, given a certain set of sensors. It is also important to point out that the environment also to some extent determines what a robot can learn about its own sensors. This was for example examined in Olsson et al. (2004a) where a robot developed its visual field in an impoverished environment with only vertical contours. Sensoritopic maps of the visual sensors showed how all sensors of the the same column were clustered together, meaning that they were completely correlated. Thus, the robot could not know, even in principle, whether these sensors were located in the same physical position or not. This work was inspired by the classical experiments by Hubel and Wiesel where kittens were reared with their visual field restricted to contours of a certain orientation (Wiesel, 1982). Their results showed how the kittens developed more neurons selective for the particular angles of contours they were restricted to while their brain developed during infancy. Unlike the kittens, the robot was able to unfold and adapt its sensoritopic map after exposure to a more "normal" environment (Olsson et al., 2004a).

The statistical structure of an agent's sensorimotor world determines not just what it can know about its world, but also how effective a certain sensor configuration is. Many changes in neuronal connections and signaling proper-

ties are due to spatial and temporal properties of incoming neuronal signals (Laughlin and Sejnowski, 2003). This has been studied in flies (Laughlin, 1981; Brenner et al., 2000) where it was shown how the encoding in contrast and motion neurons adapts to its input, similar to the entropy maximization method presented in section 2.2. In future work analysis of the statistical structure of a robot's sensory experience will be studied to understand how it can be used to optimize the sensoric system for particular environments and activities.

Acknowledgements

The work described in this paper was partially conducted within the EU Integrated Project RobotCub ("Robotic Open-architecture Technology for Cognition, Understanding, and Behaviours") and was funded by the European Commission through the E5 Unit (Cognition) of FP6-IST under Contract FP6-004370.

References

P. Andry, P. Gaussier, J. Nadel, and B. Hirsbrunner. Learning invariant sensorimotor behaviors: A developmental approach to imitation mechanisms. *Adaptive Behavior*, 12(2):117–140, 2004.

H. B. Barlow. Possible principles underlying the transformation of sensory messages. *Sensory Communication*, pages 217–234, 1961.

A Basu, I. R. Harris, and Basu S. Minimum distance estimation: The approach using density-based distances. In G. S. Maddala and C. R. Rao, editors, *Handbook of Statistics*, volume 15, pages 21–48, 1997.

L. Berthouze and Y. Kuniyoshi. Emergence and categorization of coordinated visual behavior through embodied interaction. *Machine Learning*, 31:187–200, 1998.

L. Berthouze, Y. Shigematsu, and Y. Kuniyoshi. Dynamic categorization of explorative behaviors for emergence of stable sensorimotor configurations. In *Proceedings of the International Conference on Simulation of Adaptive Behavior (SAB1998)*, pages 67–72, 1998.

D. Blank, D. Kumar, L. Meeden, and J. Marshall. Bringing up a robot: fundamental mechanisms for creating self-motivated, self-organizing architecture. *Cybernetics and Systems*, 36(2), 2005.

N. Brenner, W. Bialek, and R. de Ruyter van Steveninck. Adaptive rescaling maximizes information transformation. *Neuron*, 26:695–702, 2000.

D.M. Coppola, H. R. Purves, A. N. McCoy, and Purves D. The distribution of oriented contours in the real world. In *Proceedings of the National Academy of Sciences*, volume 95, pages 4002–4006, 1998.

T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley & Sons, Inc., N. Y., 1991.

J. P. Crutchfield. Information and its Metric. In L. Lam and H. C. Morris, editors, *Nonlinear Structures in Physical Systems – Pattern Formation, Chaos and Waves*, pages 119–130. Springer Verlag, 1990.

G. Goodhill, S. Finch, and T. Sejnowski. Quantifying neighbourhood preservation in topographic mappings. *Institute for Neural Computation Technical Report Series, No. INC-9505*, 1995.

R. Held and A. Hein. Movement produced stimulation in the development of visually guided behavior. *Journal of Comparative Physiology and Psychology*, 56:872–876, 1963.

M. I. Jordan and D. E. Rumelhart. Forward models: Supervised learning with a distal teacher. *Cognitive Science*, 16:307–354, 1992.

T. Kohonen. *Self-organizing Maps*. Springer-Verlag, Berlin, third edition, 2001.

A. Korner and H. Kraemer. Individual differences in spontaneous oral behavior in neonates. In J. Bosma, editor, *Proceedings of the 3rd Symposium on Oral Sensation and Perception*, pages 335–346, 1972.

W. J. Krzanowski. *Principles of Multivariate Analysis: A User's Perspective*. Clarendon Press, Oxford, 1988.

B. Kuipers, P. Beeson, J. Modayil, and J. Provost. Bootstrap learning of foundational representations. *Connection Science*, 18(2), 2006.

Y. Kuniyoshi, Y Yorozu, M. Inaba, and H Inoue. From visuo-motor self learning to early imitation - a neural architecture for humanoid learning. In *Proceedings of the 2003 IEEE International conference on Robotics and Automation*, pages 67–72, 2003.

M. Kuperstein. Neural model of adaptive hand-eye coordination for single postures. *Science*, 259:1308–1311, 1988.

S. B. Laughlin. A simple coding procedure enhances a neuron's information capacity. *Z. Naturforsch.*, 36c:910–912, 1981.

S. B. Laughlin and T. J. Sejnowski. Communication in neuronal networks. *Science*, pages 1870–1874, 2003.

D. D. Lee and H. S. Seung. Algorithms for non-negative matrix factorization. In *Advances in Neural Information Processing Systems 13*, pages 556–662. MIT Press, 2001.

J. Lin. Divergence measures based on the Shannon entropy. *IEEE Transactions on Information Theory*, 37(1):145–151, 1991.

B. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *Proceedings of 7th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 674–679, 1981.

M. Lungarella, G. Metta, R. Pfeifer, and G. Sandini. Developmental robotics: a survey. *Connection Science*, 15(4):151–190, 2004.

M. Lungarella, T. Pegors, D. Bulwinkle, and O. Sporns. Methods for quantifying the informational structure of sensory and motor data. *Neuroinformatics*, 3 (3):243–262, 2005.

A. Meltzoff and M. Moore. Explaining facial imitation: a theoretical model. *Early Development and Parenting*, 6:179–192, 1997.

E. A. Newman and P. H. Hartline. Integration of visual and infrared information in bimodal neurons of the rattlesnake optic tectum. *Science*, 213:789–791, 1981.

L. Olsson, C. L. Nehaniv, and D. Polani. The effects on visual information in a robot in environments with oriented contours. In *Proceedings of the Fourth International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, pages 83–88. Lund University Cognitive Studies, 2004a.

L. Olsson, C. L. Nehaniv, and D. Polani. Sensory channel grouping and structure from uninterpreted sensor data. In *2004 NASA/DoD Conference on Evolvable Hardware June 24-26, 2004 Seattle, Washington, USA*, pages 153–160. IEEE Computer Society Press, 2004b.

L. Olsson, C. L. Nehaniv, and D. Polani. Discovering motion flow by temporal-informational correlations in sensors. In *Proceedings of the Fifth International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, pages 117–120. Lund University Cognitive Studies, 2005a.

L. Olsson, C. L. Nehaniv, and D. Polani. From unknown sensors and actuators to visually guided movement. In *Proceedings of 4th IEEE International Conference on Development and Learning (ICDL 2005)*, pages 1–6. IEEE Computer Society Press, 2005b.

L. Olsson, C. L. Nehaniv, and D. Polani. Sensor adaptation and development in robots by entropy maximization of sensory data. In *Proceedings of the 6th IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA-2005)*, pages 587–592. IEEE Computer Society Press, 2005c.

L. Olsson, C. L. Nehaniv, and D. Polani. Measuring informational distances between sensors and sensor integration. In *Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems (ALIFEX)*. MIT Press, 2006. (to appear).

J.K O'Regan and A. Noe. A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24(5):1011–1031, 2001.

D Philipona, J.K. O'Regan, and J.P Nadal. Is there something out there? inferring space from sensorimotor dependencies. *Neural Computation*, 15(9), 2003.

D Philipona, J.K. O'Regan, and J.P Nadal. Perception of the structure of the physical world using multimodal unknown sensors and effectors. *Advances in Neural Information Processing Systems*, 17, 2004.

J. Piaget. *The Origins of intelligence*. Routledge, 1953.

D. Pierce and B. Kuipers. Map learning with uninterpreted sensors and effectors. *Artificial Intelligence*, 92:169–229, 1997.

D. Polani. Measures for the organization of self-organizing maps. *Self-Organizing Neural Networks: Recent Advances and Applications*, pages 13–44, 2002.

F. Rieke, D. Warland, R. de Ruyter van Stevenick, and W. Bialek. *Spikes: Exploring the neural code*. MIT Press, first edition, 1999.

P. Rochat. Self-perception and action in infancy. *Experimental Brain Research*, 123:102–109, 1998.

D. L. Ruderman. The statistics of natural images. *Network: Computation in Natural Systems*, 5:517–548, 1994.

R. Steuer, J. Kurths, C.O. Daub, Weise J., and J Selbig. The mutual information: Detecting and evaluating dependencies between variables. *Bioinformatics*, 18:231–240, 2002.

D. Stronger and P. Stone. Towards autonomous sensor and actuator model induction on a mobile robot. *Connection Science*, 18(2), 2006.

J. von Uexküll. *Streifzüge durch die Umwelten von Tieren und Menschen*. Hamburg: Rowolt, 1956.

J. Weng, J. McClelland, A. Pentland, and O. Sporns. Autonomous mental development by robots and animals. *Science*, 291:599–600, 2001.

T.N. Wiesel. Postnatal development of the visual cortex and the influence of environment. *Nature*, 299:583–591, 1982.