Proceedings Second International Symposium on New Frontiers in Human-Robot Interaction, K. Dautenhahn, J. Saunders (Eds.),
at the AISB 2010 convention, 29 March – 1 April 2010, De Montfort University, Leicester, UK

Towards Intrinsically Learned Skin Models in Robots

Simon McGregor and **Daniel Polani** [1]

Abstract. We describe the use of information-theoretic principles following [4] for sensoritopic map construction on a simulated "skin". This research constitutes a first step in robot body model learning based on uninterpreted tactile sensors. Results are reasonable, although our findings identify a serious limitation in the information distance metric: the distance is analytically bounded by the number of bins used for discretisation, so that reconstructions tend to "spherise". Furthermore, the tactile modality is significantly sparser than the visual modality, which introduces further artefacts in the information distance, particularly under maximum entropy binning. We conclude that some method which allows sensoritopic reconstruction directly from mutual information would be preferable.

1 Introduction

In many cultures, tactile contact forms an important part of human social communication. Technology permitting, the tactile modality offers opportunities for enriching and extending human-robot interaction; consequently, the development of full-surface tactile sensors (and algorithms for processing tactile data in robots) are areas of significant scientific interest and active research.

Our approach emphasises a developmental approach to robot intelligence, in which the robot's mastery of its own sensorimotor contingencies should be integral to its social behaviour. Thus, social behaviour should be built on learned abilities from the ground up, and tactile social behaviour should be based on the robot's capacity to learn about its tactile embodiment.

This research is being conducted in the context of the EU RoboSKIN project; specifically, as a precursor to tactile gesture recognition within social environments. Our aim is to provide a versatile framework for robot learning in the tactile modality, so that the robot can shape its own interactions with the physical and social world. The target platform is Kaspar, a robot play companion for children with autism, although the experiments described here are in simulation only.

A valuable first step in this direction would be a robot capable of learning its own body model automatically, based purely on information coming from its sensors. This paper describes preliminary research on tactile sensory reconstruction in 3D, extending work by Pierce & Kuipers [7] and Olsson [6, 4]. In our model, the robot's body surface is covered in tactile sensors and there is no prior information about either the shape of the robot, or the location, topology and response properties of these sensors. [7] describes this scenario as having *uninterpreted sensors and actuators*.

We consider the problem of reconstructing the geometry of a simulated passive rigid body covered in tactile sensors, with no prior information about the body or the sensors. Our experiments involve four distinct body geometries (plane, sphere, cylinder and Y-shape). No active behaviour occurs on the part of the "robot". This problem is far more basic than learning through exploration; consequently it is easier for us to conceptualise, although the inability to participate actively makes the problem "harder for the learner".

To make minimal assumptions about the sensors and the robot's environment, we use the framework of *information theory*. This framework provides a general way of conceptualising systematic relations between variables, typically at the cost of some computational complexity and data quality. A simple information-theoretic distance measure described by Crutchfield [2] serves to quantify the dissimilarity between different sensors. Using this measure and a standard multidimensional scaling algorithm [3], we can reconstruct a model loosely analogous to sensoritopic maps in the brain. The rationale for this is not necessarily to perform a geometrically faithful reconstruction of the physical placement of the sensors, but to capture their essential interrelation in some meaningful way.

The results were evaluated quantitatively by considering the linear correlation between reconstruction distance and physical sensor distance, as well as qualitatively by visual examination of the output.

2 Previous Work

The concept of learning from uninterpreted sensors and actuators is described in [7], where a variety of learning techniques are used to deduce features of the robot's sensors and actuators (such as the dimensionality of the sensor field, or the degrees of freedom in the actuators), and to learn simple control behaviours. Our group has extended information-theoretic principles to this domain, discovering visual flow information and reconstructing the visual topology of an AIBO robot (summarised in [6]).

However, previous work on sensoritopic maps from uninterpreted sensors has tended to focus on visual fields. This paper describes the application of the same information-theoretic methods to the tactile modality, which is embedded in a three-dimensional space and will typically carry sparser information than the visual modality.

3 Simulations

A number of simulations were run using the Open Dynamics Engine (ODE) physics library. Sensors were modelled as short-range proximity sensors distributed on the surface of a rigid object. These sensors contributed nothing to the dynamics of the object, and registered a signal when another simulated body came within the sensor's range. The sensor's signal was equal to the other body's maximum penetration into the sensor's sphere of detection.

We simulated four distinct sensor configurations: 60 sensors distributed uniformly on the surface of a sphere, 60 sensors distributed in 10 columns of 6 on the surface of a cylindrical object, 105 sensors distributed on the surface of a compound y-shape object, and 169

[1] University of Hertfordshire, UK, email: s.mcgregor2@herts.ac.uk

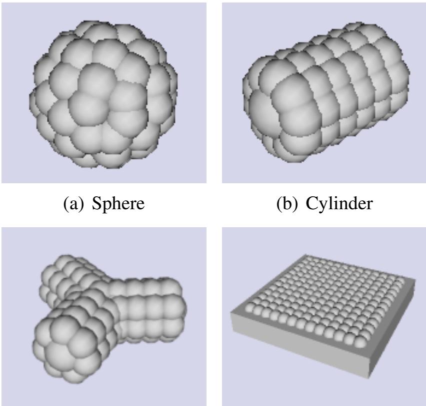

(a) Sphere (b) Cylinder

(c) Y-Shape (d) Plane

Figure 1. The sphere (60 sensors), cylinder (60 sensors, none on the ends), y-shape (105 sensors), and plane (169 sensors) geometries.

sensors distributed in a grid of 13 by 13 on one surface of a cuboid (see figure 1). The details of each geometry are given in section 3.1.

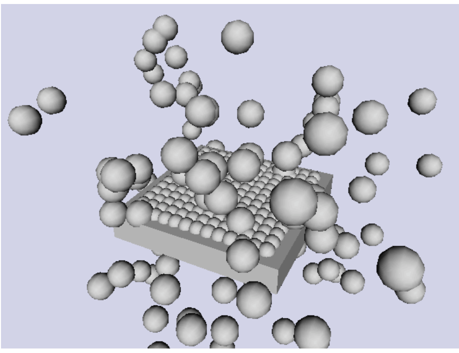

(a) Bombardment

Figure 2. Bombardment simulation scenario.

Each sensor geometry was subjected to bombardment by a "bath" of elastic balls. The simulated tactile sensor data consists simply of the response of the tactile sensors to these balls. Details of the simulation are given in section 3.2.

3.1 Geometries

3.1.1 Sphere Sensor Geometry

For the sphere, sensor positions were calculated by running a multi-dimensional scaling algorithm to embed 60 equidistant points in 3D Euclidean space. This gives a roughly even distribution over the surface of a sphere. These points were normalised to lie on the surface of a sphere of radius 0.5; the sensor range was set to 0.2m.

3.1.2 Cylinder Sensor Geometry

The cylinder object was actually modelled as a rigid cylinder capped by two semi-spheres. The cylindrical part of the object was 1m long with a radius of 0.3m (so the total length of the object was 1.6m). Sensors were placed evenly on the cylindrical portion of the object (i.e. not on the rounded ends) in 10 rows of 6 sensors, each with sensor range 0.2m. (see figure 1(c)).

3.1.3 Plane Sensor Geometry

The plane object was a 2.5m by 0.4m by 2.5m cuboid, with sensors arranged in a regular 13 by 13 planar grid on the central 2m by 2m square of the cuboid's upper surface. The sensor range was approximately 0.12m.

3.1.4 Y-Shape Sensor Geometry

This was a compound object formed by merging three rounded-end cylinders (as described in section 3.1.2): the three parts are positioned on three evenly distributed radii of a planar circle so as to resemble the letter 'Y'. Sensors were placed evenly on each cylindrical part in 7 rows of 5 sensors - note that this caused some sensors to overlap with the body of the object. The sensor range was approximately 0.21m.

3.2 Bombardment Simulation

The sensor object was an immobile body inside a large rigid cube in zero gravity. Numerous elastic spheres moved around the interior of the cube colliding with each other, the cube and the sensor object (sphere, y-shape, plane or cylinder).

The object was centred in a 4m by 4m by 4m hollow cube and was not permitted to move. 100 perfectly elastic balls of radius 0.2m and density 1kg/m^3 were initialised at random locations between 0.5m and 1.5m from the cube centre. The velocity of the balls was initialised to a uniform random vector with linear components between -100m/s and 100m/s in each direction.

The ODE simulator has a tendency to "wind down" over time, with numerical integration artefacts meaning that after a certain amount of simulation time, the mean velocity of the balls reduces substantially and eventually grinds almost to a halt. The simulation was modified to prevent this by calculating the total kinetic energy of the simulation at the start, and by every second simulation step rescaling the balls' velocity uniformly, so as to keep total kinetic energy conserved.

4 Learning The Body Model

Following [4], the algorithm we use for sensoritopic reconstruction consists of two parts:

1. Estimate the distances between each pair of sensors. These estimates are based on information distance (see section 4.1).
2. Convert these distances into a set of 3-dimensional points. This process is known as *multi-dimensional scaling* or MDS (see section 4.2).

Proceedings Second International Symposium on New Frontiers in Human-Robot Interaction, K. Dautenhahn, J. Saunders (Eds.), at the AISB 2010 convention, 29 March – 1 April 2010, De Montfort University, Leicester, UK

4.1 Information Distance

Information distance [2] is a metric defined on equivalence classes of random variables (where direct recodings of variables are equivalent). Briefly, the information distance $D(A, B)$ is equal to

$$2H(A, B) - H(A) - H(B)$$

where $H(X)$ is the Shannon entropy of X:

$$H(X) = \sum_x P(X = x) \log_2 P(X = x)$$

It can be shown [2] that information distance fulfils the metric axioms (symmetry, non-negativity, identity of indistinguishables, and the triangle inequality). This property is what allows us to use it in a standard multidimensional scaling algorithm, and why more well-known information theoretic quantities such as mutual information (a similarity measure) are not directly applicable to this scenario.

For convenience, we discretise the sensor readings into a pre-defined number of bins, which allows a straighforward computation of their individual and joint Shannon entropies according to the standard formulae. Here, we ignore the temporal sequence of the data and attend only to the co-variation in instantaneous readings. The use of information-theoretic measures captures all instantaneous relationships between the sensors, up to the resolution limits imposed by the discretisation. To get the most possible information out of the discretised sensors, we use a maximum entropy binning, which was confirmed empirically in [5] as producing better results in conjunction with the information distance metric.

4.2 Multidimensional Scaling

Multidimensional scaling (MDS) algorithms [1] are those algorithms which, given a matrix of dissimilarity scores between pairs of entities, attempt to assign those entities coordinates in a (typically) Euclidean embedding space so that the distances between the coordinate points correspond to the dissimilarity scores in the matrix. Different algorithms correspond to different ways of measuring how well the reconstructed distances approximate the dissimilarity matrix; we use the SMACOF algorithm [3], which provides a reconstruction by minimising a quantity called stress.

In our case the input to the MDS algorithm is simply the matrix of information distances between sensors, calculated as described in section 4.1. For methodological purposes, it is simplest in the context of this study to interpret this reconstruction as a direct representation of the physical locations of the sensors relative to one another. This is justified by the reasonable assumption that nearby sensors will have high mutual information and hence low information distance.

The two-fold role of the MDS algorithm in our approach should be explained. While the information distance matrix is a mathematically valid matrix of metric distances, it will typically not correspond to a set of Euclidean distances between points in 3-dimensional space. Consequently, in order to produce any coherent 3D reconstruction, the information distance matrix needs to be processed by some algorithm which performs Euclidean embedding - and this is exactly what a MDS algorithm is. However, the MDS algorithm also performs a secondary purpose. In effect, it uses our prior knowledge about the spatial structure of the problem to *improve* the distance estimate provided by the information distance matrix. We can quantify this by considering both the correlation between raw information distance and physical sensor distance, and the correlation between (post-MDS) reconstructed distance and physical sensor distance, as described in section 5.

5 Results

The original and reconstructed geometries can be seen in figures 3 to 4. It should be obvious from these images that the main features of the geometry are reconstructed for the sphere and the cylinder. When the reconstructions are viewed in three dimensions, the reconstructed plane resembles a piece of paper "scrunched up" into a ball, where local relations are reasonably preserved; this is hard to convey in a 2-dimensional image.

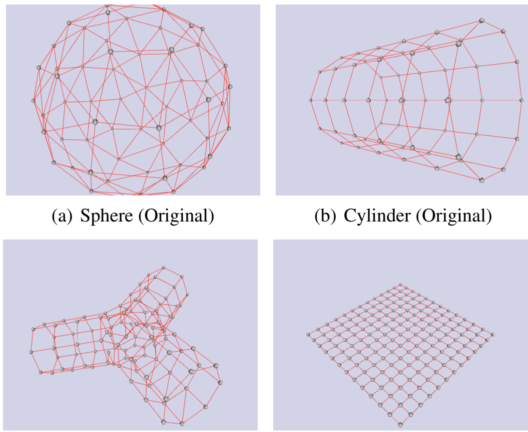

(a) Sphere (Original) (b) Cylinder (Original)

(c) Y-Shape (Original) (d) Plane (Original)

Figure 3. Original sensor geometries for sphere, cylinder, y-shape and plane.

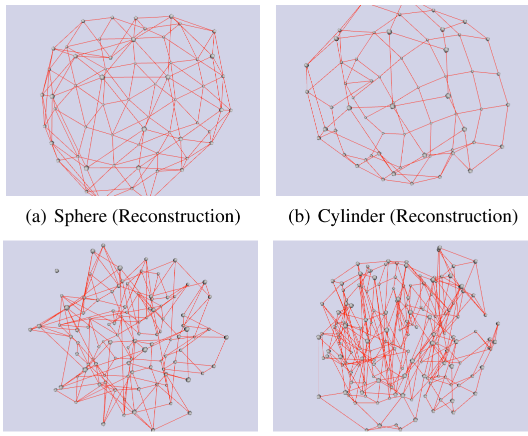

(a) Sphere (Reconstruction) (b) Cylinder (Reconstruction)

(c) Y-Shape (Reconstruction) (d) Plane (Reconstruction)

Figure 4. Sensoritopic reconstructions for sphere, cylinder, y-shape and plane.

The Y-shape is evidently less well reconstructed, and this is reflected in our correlation results (figure 5). However, it should be noted that all reconstructed distances, including these this "worse" case, have a positive correlation with the original distances which is statistically significant at the $p = 0.05$ level or better.

Form	r (pre-MDS)	r (post-MDS)
Cylinder	0.57	0.81
Y-Shape	0.24	0.34
Sphere	0.64	0.98
Plane	0.32	0.58

Figure 5. Linear correlation (Pearson's r) between simulated physical sensor distances and sensoritopic reconstruction distances before and after multi-dimensional scaling in six different experiments. The pre-MDS figures are the information distances.

It can be seen that, in most cases, application of the multi-dimensional scaling algorithm makes a dramatic difference to the quality of the reconstruction. This should not be surprising, since MDS imposes a Euclidean structure on the reconstruction which is not necessarily there in the information distance.

6 Discussion

Our results using this method are generally weaker than the results obtained by Olsson [4]. This has allowed us to analyse the limitations of the technique and identify important differences between the tactile and visual modalities.

The technique's biggest limitation relates to a hard upper bound on the information distance metric, which causes the majority of reconstructed inter-point distances to saturate at a maximum value. The upper bound arises as follows: the information metric $D(A, B)$ can be written as

$$H(A) + H(B) - 2I(A; B)$$

where $H(X)$ is the Shannon entropy of the random variable X and $I(X; Y)$ is the mutual information of X and Y. Since all these quantities are non-negative, the largest value $D(A, B)$ can take is $H(A)+H(B)$. But our discretisation into n bins means that the maximum entropy for any sensor is simply $\log_2 n$, regardless of the nature of the sensor. Moreover, our maximum entropy binning method means that $H(A) + H(B)$ is approximately $2 \log_2 n$ for most sensor pairs.

This constant-distance phenomenon causes the reconstructions to "warp" such that the reconstructed sensor positions tend to be distributed on the surface of a sphere (see figures 4, although the phenomenon is more evident when viewing in three dimensions). This phenomenon is in fact also visible in Olsson's reconstructions (see, e.g. figure 2c in [4]). However, our scenario is different from Olsson's in several ways which we believe significantly exacerbate the problem, reducing the effectiveness of the information distance method:

- Our reconstruction is in 3D, while Olsson's was in 2D. Consequently there is more hypersphere surface available for sensor points to occupy and non-spherical shapes will experience more distortion.
- Our data is sparse, which reduces the correlation between information distance and physical distance. The maximum entropy binning divides data as evenly as possible amongst bins; if the data is rich, the entropy of each sensor will be the same (equal to the log of the number of bins). In practice, then, only the mutual information between sensors is reflected in the information distance. When data is more than 50% sparse, however, the bin containing 0 will be larger than the other bins, and the entropy of each individual sensor accounts for significant variation in the information distances. This acts as noise, and mutual information is the better indicator of physical distance.

The direct use of mutual information (as opposed to information distance) promises more accurate results. However, mutual information is a similarity measure rather than a dissimilarity measure, so it is not appropriate as an input to classical MDS techniques. Our most recent research focuses on ways to overcome this problem.

ACKNOWLEDGEMENTS

The work described in this paper is conducted within the EU project RoboSKIN and is funded by the European Commission in the 7th Framework Program under grant agreement 231500.

REFERENCES

[1] I. Borg and P.J.F. Groenen, *Modern Multidimensional Scaling: Theory and Applications*, Springer, 2005.
[2] J.P. Crutchfield, 'Information and its metric', in *Nonlinear Structures in Physical Systems - Pattern Formation, Chaos and Waves*, pp. 119–130. Springer Verlag, (1990).
[3] Jan de Leeuw, 'Applications of convex analysis to multidimensional scaling', in *Recent Developments in Statistics*, eds., J.R. Barra, F. Brodeau, G. Romier, and B. Van Cutsem, 133–146, North Holland Publishing Company, Amsterdam, (1977).
[4] Lars Olsson, Chrystopher L. Nehaniv, and Daniel Polani, 'Measuring informational distances between sensors and sensor integration', in *In Artificial Life X - Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, eds., Luis Mateus Rocha, Larry S. Yaeger, Mark A. Bedau, Dario Floreano, Robert L. Goldstone, and Alessandro Vespignani, pp. 316–323. MIT Press, (2005).
[5] Lars Olsson, Chrystopher L. Nehaniv, and Daniel Polani, 'Sensor adaptation and development in robots by entropy maximization of sensory data', in *In Proceedings of the 6th IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA 2005). IEEE Computer*, pp. 587–592. Society Press, (2005).
[6] Lars Olsson, Chrystopher L. Nehaniv, and Daniel Polani, 'From unknown sensors and actuators to actions grounded in sensorimotor perceptions', *Connect. Sci.*, **18**(2), 121–144, (2006).
[7] David Pierce and Benjamin Kuipers, 'Map learning with uninterpreted sensors and effectors', *ARTIFICIAL INTELLIGENCE*, **92**, 169–229, (1997).