

Information-theoretic approach to embodied category learning

Gabriel Gómez¹, Max Lungarella² and Danesh Tarapore³

¹*Artificial Intelligence Laboratory, University of Zurich, Switzerland*

²*Dept. of Mechano-Informatics, the University of Tokyo*

³*Autonomous Systems Lab., EPFL, Lausanne, Switzerland*

gomez@ifi.unizh.ch, maxl@isi.imi.i.u-tokyo.ac.jp, danesh.tarapore@epfl.ch

Abstract

We address the issue of how statistical and information-theoretic measures can be employed to quantify the categorization process of a simulated robotic agent interacting with its local environment. We show how correlation, entropy, and mutual information can help identify distinct informational structure which can be used for object classification. Further, by means of the isometric feature mapping algorithm, we analyze the weights of a neural network designed to find clusters based on these distinct information theoretic characteristics of the object’s shape, size and color. We conclude that an understanding of the information-theoretic implications of categorization could help design robots with improved categorization and better exploration strategies.

1 introduction

In our daily lives, we are exposed to a barrage of multimodal sensory stimulation (e.g., visual, tactile, proprioceptive). Perceptual categorization can be conceptualized as the ability to identify regularities in this continuously changing stream of sensory information, and to treat similar, but not necessarily identical objects and events, as being somehow equivalent. Recently, evidence has been accumulating showing that perceptual categorization is not a mere mapping from a set of sensory nodes to a set of category nodes – as previously assumed – but is instead the result of a process of sensorimotor coordinated interaction between an embodied agent and its local environment. Such interaction has been hypothesized to be one of the major information theoretic implications of embodiment, because it allows an agent to actively generate constraints in its sensory input ([1, 2, 3, 4]). Such constraints, in turn, simplify the problem of learning categories by inducing spatio-temporal correlations, and by reducing drastically – in an information theoretic

sense – the number of degrees of freedom (dimensionality) of the sensory space.

This paper aims at quantitatively underpinning this hypothesis (see also [1, 2, 5, 6, 7, 8]). We present the analysis of data collected by a simulated robotic agent by means of correlation, mutual information, geometric separability index, and isometric feature maps. We conclude that an information-theoretic approach to the study of categorization leads to a better (quantitative) understanding of how embodied interaction simplifies category learning despite the high dimensionality of the sensorimotor data sets. This approach may also shed light on the characteristics of the objects being categorized, and reduce the time required by the agent to learn.

2 Experimental Setup

We simulated a two-wheeled robot moving in a closed environment cluttered with randomly distributed, colored objects. Objects in the environment were red cubes and cylinders.

The robot was equipped with 11 proximity sensors (d_{0-10}) and a pan-controlled image sensor or camera unit (see Fig. 1b). The proximity sensors had a position-dependent range (see caption of Fig. 1). The output of each sensor was affected by additive white noise, and was partitioned into a space having 32 discrete states, leading to sensory signals with a 5 bit resolution. To reduce the dimensionality of the input data, we divided the camera image into 24 vertical rectangular slices (i_{0-23}). We computed the amount of “effective” red color in each slice as $R=r/(b+g)/2$.

For the control of the robot we opted for the Extended Braitenberg Architecture [3] (see Fig. 2). To pick up the regularities induced by the sensorimotor coordination we used a Kohonen feature map [9].

Before being projected onto the Kohonen map, the 28-dimensional input vector consisting of the ac-




Figure 1: (a) Bird's eye view of the robot and its ecological niche. The trace represents a typical path of the robot during an experiment. (b) Schematic representation of the agent. The distance sensors have a range that depends on their position. If rl is the length of the robot, then the range of d_0 , d_1 , d_9 , and d_{10} is $1.8rl$, the one of d_2 and d_3 is $1.2rl$, and the one of d_4 , d_5 , d_6 , d_7 , and d_8 is $0.6rl$.




Figure 2: Block diagram of the control architecture.

tivations of 24 red channels, the activations of two out of four proximity sensors located on the same side of the robot (that is, either the pair d_2, d_9 , or the pair d_3, d_{10}), as well as the left and right motor activation values (m_l, m_r), was preprocessed by the “input selector” (see Fig. 2). If the agent circled an object counter-clockwise – i.e., the object was on its left – the input vector of the Kohonen map was $[i_0, i_1, \dots, i_{23}, d_2, d_9, m_l, m_r]^T$. However, if the agent circled an object clockwise – i.e., the object was on its right – the input vector was $[i_{23}, i_{22}, \dots, i_0, d_3, d_{10}, m_r, m_l]^T$. That is, the vector (i_0, \dots, i_{23}) was fed to the Kohonen map in reverse order (i_{23} was the first element, followed by d_3 and d_{10} ,

and the right and left motor activation). The reason for this re-ordering operation was to avoid having the Kohonen map discriminate between objects, based on the direction of circling.

The 28-dimensional output vector of the input selector was normalized to unit length, and each of its elements was projected onto each of the neurons of the Kohonen map (that is, the input layer was fully connected to the map layer). The Kohonen map consisted of 576 neurons, arranged in the form of two-dimensional lattice with $N_r=24$ rows and $N_c=24$ columns. The network's $24 \times 24 \times 28=16128$ initial synaptic weights were chosen from a random set. The dependence of the learning-rate parameter η on discrete time n was chosen to be $\eta(n) = \eta_0 e^{-n/\tau_1}$, where $\eta_0=1.0$ was the initial value of the learning parameter, and $\tau_1 = 2.2^{\log N_r N_c}$. That is, the learning rate decreased exponentially over time. Another feature of the Kohonen map was that the size of the neighborhood of each neuron shrank over time. The dependence on discrete time n was $\sigma(n) = \sigma_0 e^{-n/\tau_2}$, where $\sigma_0=9.0$ and $\tau_2 = 2.0^{\log N_r N_c}$.

3 Behavioral specification

At the outset of each experimental run the behavioral mode of the robot was set to “exploring.” In this mode the robot roams through the environment at a constant speed while avoiding obstacles. Upon detection of a salient red object, the robot approaches it under guidance of the visual system (“tracking” mode). As soon as the object is close, the robot starts circling around it keeping it in the center of its visual field by adjusting the camera's pan-angle (“circling” mode). Concurrently, a habituation signal starts increasing. While the robot circles the object, the values of the input signal (24 image sensors, a pair of side distance sensors, and two motor activations) – after being appropriately re-organized by the input selector – are projected onto the Kohonen feature map; the synaptic weights of the network are updated, and the learning-rate parameter decreased. The robot keeps circling around the object for a while and then resumes the exploration of its environment. The trace of a typical experimental run is shown in Fig. 1a. The entire experiment comprised 78 such experimental runs, each run consisting of approximately 6000 data samples. All samples were stored into a time series file for subsequent analysis.

The categorization error as the percentage of misclassified objects per unit of time is shown in Fig. 3. Clearly, the network learned to discriminate between




Figure 3: Categorization error (vertical axis) versus time (horizontal axis). (a) cubes; (b) cylinders.

cubes and cylinders: The classification error after timestep 77 dropped to almost zero.

4 Methods

In this section, we describe the statistical and information theoretical measures employed in this paper. The correlation $Corr(X, Y)$ quantifies the amount of linear dependency between two random variables X and Y (e.g., two sensory channels), and is given by $\sum_{x \in X} \sum_{y \in Y} p(x, y) (x - m_X)(y - m_Y) / \sigma_X \sigma_Y$, where $p(x, y)$ is the second order (or joint) probability density function, m_X and m_Y are the mean, and σ_X and σ_Y are the standard deviation of x and y computed over X and Y . The entropy of a random variable X is a measure of its uncertainty, and is defined as $H(X) = -\sum_{x \in X} p(x) \log p(x)$, where $p(x)$ is the first order probability density function associated with X (in a sense entropy provides a measure for the sharpness of $p(x)$). Similarly, the joint entropy between variables X and Y is defined as $H(X, Y) = -\sum_{x \in X} \sum_{y \in Y} p(x, y) \log p(x, y)$. For entropy as well as mutual information, we assumed the binary logarithm. Using the joint entropy $H(X, Y)$, we can define the mutual information between X and Y as $MI(X, Y) = H(X) + H(Y) - H(X, Y)$.

To get a better grasp on the regularities in the sensory space, we computed also the Geometric Separability Index (GSI) introduced in [10]. Geometric separability is a generalization of linear separability, and quantifies how close regions in the sensory space belonging to the same object are to each other. The GSI is computed by checking for every sensory pattern whether the nearest pattern (in terms of the Euclidean distance) is part of the same class. It is calculated as

follows:

$$\sum_{i=1}^n \frac{(f[x_i] + f[x_i^*] + 1) \bmod 2}{N},$$

where $f[x_i] = 0$, if x_i belongs to one class, and $f[x_i] = 1$, if x_i belongs to the other class (f is also called the category function); x_i is the i -th sensory pattern (N -dimensional vector), and x_i^* is the nearest neighbor of x_i . If the nearest pattern in sensory space always belong to the same class of the currently perceived object the GSI is 1. High values of the GSI thus indicate that the sensory patterns belonging to the two categories are quite separated in the input space and easy to discriminate, while value close to 0.5 indicate that the sensory patterns corresponding to the two categories completely overlap.

In addition to these three measures, we also used the isometric feature mapping, or Isomap, algorithm described in [11]. Isomap solves the problem of dimensionality reduction by using local metric information to learn the underlying global geometry of a data set. It discovers – given only the unordered high-dimensional input – the low-dimensional representations of the data. Although the classical techniques for dimensionality reduction such as Principal Component Analysis (PCA) and Multi Dimensional Scaling (MDS) are simple to implement and efficiently computable, many data sets contain non linear structures that are invisible to PCA and MDS. Isomap combines the major algorithmic features of PCA and MDS with the flexibility to discover nonlinearity in the input data.

5 Data Analysis and Results

Correlation

While the robot circled the object, we observed negative correlation between the left and right motors. The circling behavior also induced negative correlation between pairs of proximity sensors located on the same side of the robot, that is, between (d_2, d_9) , (d_0, d_2) and (d_2, d_4) when the robot circled the object in the counter-clockwise direction. (see Fig. 4).

There is a high discrepancy in the (d_0, d_2) sensor correlation between cubes (-0.3083) and cylinders (0.7109). In the case of cylinders, due to their smooth surface, the sensors d_2 as well as d_0 will always be in contact with the object. Due to the corners of the cubes, as well as the length of its edge, this positive correlation will not be evident.

Correlation	Cubes	Cylinders
(d_2, d_9)	-0.4995	-0.3189
(d_0, d_2)	-0.3083	0.7109
(d_2, d_4)	-0.4476	0.0540

Table 1:




Figure 4: Distance sensor correlation matrix obtained from the pair-wise correlation between pairs of distance sensors (indexes d_0 to d_{10}) for (a) cubes and (b) cylinders. The behavioral state is “circling”. The higher the correlation the larger the size of the square.

While the robot was circling around a cube, strong correlations between the output of all the 24 red channels (i_1 to i_{24}) could be observed. The average correlation amounted to 0.5628 (see Fig. 5a). For cylindrical objects, the average correlation between the same red channels (i_1 to i_{24}) was 0.1321.




Figure 5: Image sensor correlation matrix obtained from the pair-wise correlation between pairs of image sensors (indexes 1 to 24) for (a) cubes and (b) cylinders. The behavioral state is “circling”. The higher the correlation the larger the size of the square.

Entropy and mutual information

The pair-wise mutual information between the 24 image sensors, and the 11 proximity sensors is shown

in Fig. 6 and Fig. 7. The diagonals of the plots represent the entropy of the sensory stimulation.

While the robot circled the object, we observed high mutual information between pairs of proximity sensors located on the same side of the robot, that is, between (d_2, d_9) , (d_0, d_2) and (d_2, d_4) when the robot circled the object in the counter-clockwise direction.

Mutual Information	Cubes (bits)	Cylinders (bits)
(d_0, d_0)	0.0464	0.1752
(d_2, d_2)	3.0225	3.0439
(d_9, d_9)	1.0292	0.9103
(d_0, d_2)	0.0252	0.1475
(d_2, d_4)	0.1956	0.0506
(d_2, d_9)	0.1919	0.0858

Table 2:

There is a high discrepancy in the sensor d_0 entropy between cubes (0.0464) and cylinders (0.1752). This discrepancy is also observed in (d_0, d_2) mutual information between cubes (0.0252) and cylinders (0.1474). In the case of cylinders, due to their smooth round surface, the sensor d_0 will always be in contact with the object. However the sharp corners of the cube prevent this contact from occurring as frequently. We also observe a high discrepancy in (d_2, d_4) mutual information between cubes (0.1956) and cylinders (0.0506).




Figure 6: Distance sensor mutual information matrix obtained from mutual information between pairs of distance sensors (d_0 to d_{10}), in one particular experimental run. The behavioral state is “circling” (a) cubes and (b) cylinders. The higher the mutual information the larger the size of the square.

While the robot was circling around a cube, high mutual information between the output of all the 24 red channels (i_1 to i_{24}) could be observed. The average mutual information amounted to 0.8166 bits (see Fig. 7a). For cylindrical objects, the average mutual information between the same red channels (i_1 to i_{24}) was 0.2676 bits.




Figure 7: Image sensor mutual information matrix obtained from the pair-wise mutual information between pairs of image sensors (indexes 1 to 24), in one particular experimental run. The behavioral state is “circling” (a) cubes and (b) cylinders. The higher the mutual information the larger the size of the square.

Geometric Separability Index (GSI)

The GSI identifies the regularities founded by the Kohonen map in the sensory space, allowing us to predict the Kohonen map’s learning performance. A GSI of 0.9989 when the robot has learnt to classify the cubes and cylinders. Sensory motor interactions indicate the presence of clusters in the data, which the Kohonen map has extracted. Indeed, the “circling” behavioral mode by inducing constraints in the input space, made cubes and cylinders discernible because data points belonging to one object were close to each other in the data space (correlation). The data points belonging to different objects are far apart.




Figure 8: Dimensionality reduction of the 28-dimensional sensory input vector. Applied to $N=576$ vectors with $K=7$, Isomap learned a two-dimensional embedding of the data’s informational structure.

Isometric feature mapping

To get a grasp on the high-dimensional space spanned by the sensory data, we used the Isomap algorithm. By applying Isomap to the 28-dimensional input space, we were able to compress it to a two-dimensional space. In Fig. 8, the clusters have been enclosed in boxes. The X axis shows positive correlation with a linear combination of the following factors: a) the difference between distance sensors d_2-d_9 or d_3-d_{10} : cubes are clustered by higher values of this difference and cylinders by the lower values; b) the difference of right and left motor m_R-m_L : the cubes are clustered according to higher values of this difference compared to the cylinders; and c) the sum of the image sensor values focused on the object being circled: cubes display a larger number of image sensor activations than cylinders. The total correlation considering all these factors was 0.931.

Activation of the sensor d_2 orients the robot toward objects (attraction); activation in d_9 , on the other hand, makes the robot turn away from objects (repulsion). For cubes, due to their flat surfaces and sharp corners, attraction prevails on repulsion. This results in cubes being associated with larger values of the differences d_2-d_9 , and d_3-d_{10} . In the case of cylinders, however, attraction prevails on repulsion, causing the cylinders to be associated with low values of the difference d_2-d_9 or d_3-d_{10} . Due to the sudden turns the robot has to make when circling cubes, the difference m_R-m_L is higher for cubes than for cylinders. When circling cubes, the robot is closer to it than for cylinders. The reason behind this is that in the case of cylinders, the robot starts circling the object, when one of its distance sensors d_4 or d_5 detects the object ahead. However, in the case of cubes, the distance sensors d_4 or d_5 do not detect the cube. Rather, it is detected by the center distance sensor d_6 which has a much shorter range. This causes the robot to be much closer to the encircled object in the case of cubes (larger number of image sensor activations) than in the case of cylinders (smaller number of image sensor activations).

6 Further discussion and conclusion

Does sensorimotor coordinated activity induce discernible regularities or structure in the sensorimotor space? As shown by the statistical and information-theoretic analyses of the recorded sensory and motor data, appropriate coordinated motor activity leads to

a characteristic “fingerprint” of the sensorimotor interaction – that is, spatio-temporal patterns in the sensorimotor space reproducible across multiple experimental runs (for a similar result see also [7]). Such patterns can be identified, and exploited to simplify subsequent discrimination tasks.

It is important to note that the categorization for the robot is based on a 28-dimensional vector of sensorimotor values, where each of the sensors take 32 values (each sensor has a 5 bit resolution). Indeed, this results in a very large space of potential sensorimotor configurations (28^{32}). The circling behavior of the robot, allows the agent to generate constraints in its sensory input that lead to a drastic dimensionality reduction of the sensory space. This greatly reduces the search space of the robot. In other words, the circular behavior, akin to the object rotation behavior found in human infants [12], or exploration strategies observed in adults exploring objects with their hands [13], induces spatio-temporal correlations among the sensory patterns. These correlations are a further indication of the fact that the sensorimotor coupling leads to a reduction of the degrees of freedom in the input space [2].

The high values of correlation that we have observed in the “circling” behavioral state allow us to infer that while in this state, the agent generates redundancy in the sensory signals, which can be picked up by the agent itself. In turn, such redundancy is a prerequisite for learning, and provides a basis for some sort of “self-cognition.” That is, the agent can acquire a notion of its own emergent behavior.

We further note that while it is possible, at least in simple cases, to characterize exploration strategies by means of information theoretic and statistical measures, the proposed measures are by no means equivalent, but complement each other. For instance, in contrast to linear correlation, mutual information takes into account also nonlinear dependencies between various stochastic variables. The geometric separability index, on the other side, discerns sensory input based on how linearly separable it is, that is, based on where it falls respect to a hyperplane. Isomap, finally, tries to isolate nonlinear degrees of freedom in the data set. It would be of interest to understand if the agent could perform such “analyses” on its own. By characterizing its interactions with the environment, it would form the basis on which to build its own individual experiences, its memory. In future work, we also plan to study more complex agent-object interactions, and to experiment with more sophisticated robots.

Acknowledgements

Gabriel Gómez would like to thank the EU-Project ADAPT (IST-2001-37173) for funding. Max Lungarella would like to thank the University of Tokyo for funding and Yasuo Kuniyoshi for support and encouragement.

References

- [1] R. Pfeifer and C. Scheier. Sensory-motor coordination: The metaphor and beyond. *Robotics and Autonomous Systems*, 20:157–178, 1997.
- [2] C. Scheier, R. Pfeifer, and Y. Kuniyoshi. Embedded neural networks: Exploiting constraints. *Neural Networks*, 11:1551–1569., 1998.
- [3] R. Pfeifer and C. Scheier. *Understanding Intelligence*. MIT Press, 1999.
- [4] M. Lungarella, T. Pegors, D. Bulwinkle, and O. Sporns. Methods for quantifying the informational structure of sensory and motor data. *Neuroinformatics (submitted)*, 2004.
- [5] M. Lungarella and R. Pfeifer. Robots as cognitive tools: information-theoretic analysis of sensory-motor data. In *Proc. of 2nd Int. IEEE/RSJ Conf. on Humanoid Robotics*, pages 245–252, 2001.
- [6] O. Sporns. Generating structure in sensory data through coordinated motor activity. In *Proc. of Joint Conf. on Neural Network*, page 2796, 2003.
- [7] D. Tarapore, M. Lungarella, and G. Gómez. Quantifying patterns of agent-environment interaction. *Robotics and Autonomous Systems (accepted for publication)*, 2005.
- [8] R. Te Boekhorst, M. Lungarella, and R. Pfeifer. Dimensionality reduction through sensory-motor coordination. In *Proc. of the Joint Int. Conf. ICANN/ICONIP*, pages 496–503, 2003.
- [9] T. Kohonen. *Self-organization and associative memory*. Springer, Berlin, 1988.
- [10] C. Thornton. Separability is a learners best friend. In *Proc. 4th Neural Computation and Psychology Workshop: Connectionist representations*, pages 40–47, 1997.
- [11] J. Tenenbaum, V. de Silva, and J. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290:2319–2323, 2000.
- [12] H. A. Ruff. Infants’ manipulative exploration of objects: effects of age and object characteristics. *Developmental Psychology*, 20:9–20, 1984.
- [13] S. J. Lederman and R. L. Klatzky. Haptic exploration and object representation. In M. Goodale, editor, *Vision and Action: The Control of Grasping*, pages 98–109. New Jersey: Ablex, 1990.