ADAPT
IST-2001-37173
Artificial Development Approach to Presence Technologies

Deliverable Item 3.3
A set of formal methods for the analysis of the interplay of morphology, materials and control

Delivery Date: July 10, 2005
Classification: Public
Responsible Person: Gabriel Gómez
Partners Contributed: ALL

Short Description: One of the fundamental methodological problems of studying behavior in the real world is that the various components of an agent (i.e., sensory, motor and neural systems) are strongly dependent on one another and on the specific interactions with the environment. The synthetic methodology serves to extract and keep track of these dependencies: the sensory-motor and internal states of the robot can be recorded into a time-series file and subsequently analyzed.
We investigate by means of statistical and information-theoretic measures, to what extent sensory-motor coordinated activity can generate and structure information in the sensory channels of an agent interacting with its surrounding environment. The results show how the usage of correlation, entropy, and mutual information can be employed (a) to segment an observed behavior into distinct behavioral states, (b) to analyze the informational relationship between the different components of the sensory-motor apparatus, and (c) to identify patterns (or fingerprints) in the sensory-motor interaction between the agent and its local environment.

Project funded by the European Community under the "Information Society Technologies" Programme (1998–2002)

1 Introduction

Manual haptic perception is the ability to gather information about objects by using the hands. Haptic exploration is a task-dependent activity, and when people seek information about a particular object property, such as size, temperature, hardness, or texture, they perform stereotyped exploratory hand movements. In fact, spontaneously executed hand movements are the best ones to use, in the sense that they maximize the availability of relevant sensory information gained by haptic exploration (Lederman and Klatzky, 1990). The same holds for visual exploration. Eye movements, for instance, depend on the perceptual judgement that people are asked to make, and the eyes are typically directed toward areas of a visual scene or an image that deliver useful and essential perceptual information (Yarbus, 1967). To reason about the organization of saccadic eye movements, Lee and Yu (1999) proposed a theoretical framework based on information maximization. The basic assumption of their theory is that due to the small size of our foveas (high resolution part of the eye), our eyes have to continuously move to maximize the information intake from the world. Differences between tasks obviously influence the statistics of visual and tactile inputs, as well as the way people acquire information for object discrimination, recognition, and categorization.

Clearly, the common denominator underlying our perceptual abilities seems to be a process of sensory-motor coordination which couples perception and action. It follows that coordinated movements must be considered part of the perceptual system (Thelen and Smith, 1994), and whether the sensory stimulation is visual, tactile, or auditory, perception always includes associated movements of eyes, hands, arms, head and neck (Ballard, 1991; Gibson, 1988). Sensory-motor coordination is important, because (a) it induces correlations between various sensory modalities (such as vision and haptics) that can be exploited to form cross-modal associations, and (b) it generates structure in the sensory data that facilitates the subsequent processing of those data (Lungarella and Pfeifer, 2001; Lungarella and Sporns, 2004; Nolfi, 2002; Sporns and Pegors, 2003).

Our long-term goal is to quantitatively understand what sort of coordinated motor activities lead to what sort of information. We also aim at identifying "fingerprints" (or patterns of sensory or sensory-motor activation) characterizing the agent-environment interaction. Our approach builds on top of previous studies on category learning (Pfeifer and Scheier, 1997; Scheier and Pfeifer, 1997), as well as on work on the information-theoretic and statistical analysis of sensory and motor data (Lungarella and Pfeifer, 2001; Sporns and Pegors, 2003; Te Boekhorst et al., 2003).

2 Methods

First, we introduce some notation. Correlation quantifies the amount of linear dependency between two random variables X and Y, and is given by the following formula:

$$Corr(X,Y) = (\sum_{x \in X} \sum_{y \in Y} p(x,y)\,(x-m_X)(y-m_Y))/\sigma_X\,\sigma_Y \qquad (1)$$

where:
$p(x,y)$ is the second order (or joint) probability density function,

m_X and m_Y are the means,
and σ_X and σ_Y are the standard deviations of x and y computed over X and Y

Note that the analyses were performed by fixing the time lag between the two time series to zero.

The entropy of a random variable X is a measure of its uncertainty, and is defined as:

$$H(X) = -\sum_{x \in X} p(x) \, log \, p(x) \tag{2}$$

where:
$p(x)$ is the first order probability density function associated with X,

In a sense entropy provides a measure for the sharpness of $p(x)$.

The joint entropy between variables X and Y is defined analogously as:

$$H(X,Y) = -\sum_{x \in X} \sum_{y \in Y} p(x,y) \, log \, p(x,y) \tag{3}$$

Mutual information measures the statistical independence of two random variables X and Y (Cover and Thomas, 1991; Shannon, 1948). Using the joint entropy $H(X,Y)$, we can define the mutual information between X and Y as:

$$MI(X,Y) = H(X) + H(Y) - H(X,Y) \tag{4}$$

In comparison with correlation, mutual information provides a better and more general criterion to investigate statistical dependencies between random variables (Steuer et al., 2002). For entropy as well as for mutual information, we assumed the binary logarithm.

Correlation, entropy and joint entropy were computed by first approximating $p(x)$ and $p(x,y)$. The most straightforward approach is to use a histogram-based technique, described, for instance, in (Steuer et al., 2002). Because the sensors had a resolution of 5 bits, we estimated the histograms by setting the number of bins to 32 (which led to a bin-size of one). Having a unitary bin size allowed us to map the discretized value of the sensory stimulus directly onto the corresponding bin for the approximation of the joint probability density function. Because of the limited number of available data samples, the estimates of the entropy and of the mutual information were affected by a systematic error (Roulston, 1999). We compensated for this bias by adding a small corrective term T to the computed estimates: $T = (B-1)/2\,N$ to the entropy estimate (where N is the size of the temporal window over which the entropy is computed, and B is the number of states for which $p(x_i) \neq 0$), and $T = (B_x + B_y - B_{x,y} - 1)/2\,N$ to the mutual information estimate (where B_x, B_y, $B_{x,y}$, and N have an analogous meaning to the previous case).

3 Experimental Setup

We have performed experiments with three different experimental setups:

3.1 hand-eye coordination setup

Manipulation allows a robot to categorize different objects, but first it has to learn to manipulate those objects, bringing them closer to its "fovea" for further and detailed exploration. For this we used a 6 DOF robot arm and an active vision system with 4 DOF (see Figure 1), which task was to learn how to move a colored object on its griper from the periphery to the center of its visual field ("fovea") (Gómez and Eggenberger Hotz, 2004a,b).

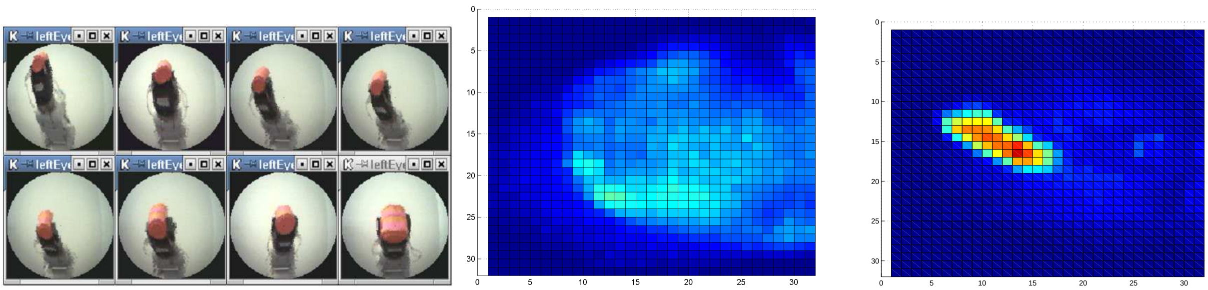

(a) (b) (c)

Figure 1: (a) Robot's perspective of the position of an object during the learning. (b) Total amount of red during random exploration. (c) Total amount of red during a sensory motor coordinated interaction.

A new set of experiments is being performed; Figure 2 shows the final position of the robot arm with different objects on its gripper which are then rotated in front of the cameras for a little while.

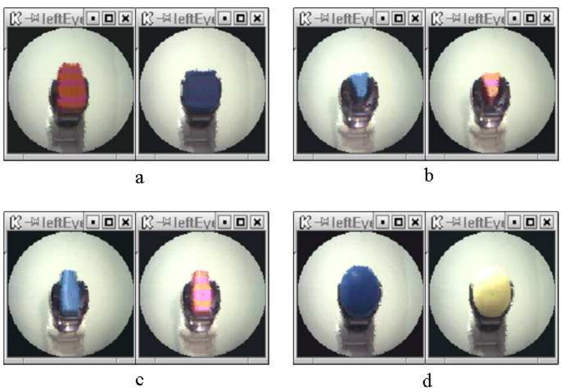

Figure 2: The robot using vision and touch explores objects of different colors, shapes and materials. (a) cube, (b) pyramid, (c) cylinder, and (d) ball.

3.2 Antropomorphic robot hand setup

Figure 3 shows a 13 degrees of freedom robotic hand equipped with bending and pressure sensors whose task was to learn to grasp objects regardless of their shape based on tactile information (Gomez et al., 2005).

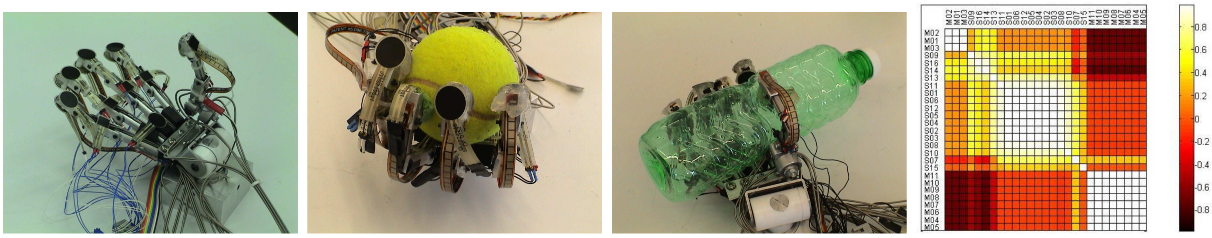

(a) (b) (c) (d)

Figure 3: (a) Tendon driven robotic hand. (b and c) Grasping different objects. (d) Correlation matrix obtained from the pair-wise correlation of the bending sensors, pressure sensors, and motors for one particular experimental run.

3.3 simulated mobile robotic-agent setup

A simulated two-wheeled robot wondering around in a closed environment cluttered with randomly distributed, colored cylindrical objects. Figure 4 a shows the echological niche of the robot. The robot was equipped with eleven proximity sensors (d_{0-10}) and a pan-controlled camera unit (see Fig. 4 b). The proximity sensors had a position-dependent range, that is, the sensors in the front and the ones in the back had a short range, whereas the ones on the sides had a longer range. The output of each sensor was affected by additive white noise with an amplitude of 10% the sensor range, and was partitioned into a space having 32 discrete states, leading to sensory signals with a 5 bits resolution. To reduce the dimensionality of the input data, we divided the camera image into 24 vertical rectangular slices (i_{1-24}), with width of two pixels for the slices close to the center (i_{7-18}), and width of six pixels for the slices in the periphery (i_{1-6} and i_{19-24}). We computed the amount of the "effective" red color in each slice as $R = r - (b + g)/2$, where r, g, and b are the red, green, and blue components of the color associated with each pixel of the slice. The control of the robot was based on the Extended Braitenberg Architecture (Pfeifer and Scheier, 1999).

4 Experiments

In this section we concentrate in the experiments performed using the simulated agent described above, but the same kind of analysis is being performed for the other robotic setups.

Because our goal is to illustrate how standard statistical and information-theoretic measures can be employed to quantify (and fingerprint) the agent-environment interaction, we started by decomposing the robot's behavior into three distinct behavioral states: (a) "explore the environment" and "find red objects", (b) "track red objects", and (c) "circle around red objects." A top view of a typical experiment is shown in Fig. 4 a. At the outset of each experimental run, the robot's initial position was set to the final position of the previous experiment (except

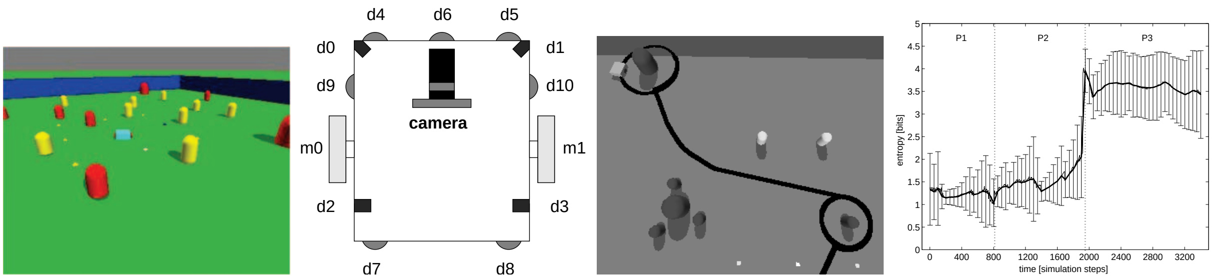

(a) (b) (c) (d)

Figure 4: (a) Bird's eye view on the robot and its ecological niche. (b) Schematic representation of the agent. (c) The trace represents a typical path of the robot during one experiment, three behavioral states are identified (exploring, tracking and circling). (d) Entropy of the effective red color averaged over all vertical slices. P1: exploring; P2: tracking; P3: circling.

for the first experiment where the robot was placed in the origin of the x-y plane), and the behavioral state was reset to "exploring." In this particular state, the robot randomly explored its environment while avoiding obstacles. Concurrently, the robot's camera panned from side to side (by 60 degrees on each side). If the maximum of the effective red color (summed over the entire image) passed a given (fixed) threshold, it was assumed that the robot had successfully identified a red object. The behavioral state was set to "tracking", the camera stopped rotating from side to side, and the robot started moving in the direction pointed at by the camera, trying to keep the object in the camera's center of view. Once close to the red object, the robot started circling around it (while still keeping it in its center of view by adjusting the camera's pan-angle). At the same time, a "boredom" signal started increasing. The robot kept circling around the object, until the boredom signal crossed an upper threshold. In that instant, the robot stopped circling, and started backing away from the red object, while avoiding other objects. Concurrently, the boredom signal began to decrease. When the boredom signal finally dropped below a lower threshold, the robot resumed the exploration of the surrounding environment. We performed 16 experiments, each of which lasted approximately 3400 time steps. The sensor and motor activations were stored into a time series file for subsequent analysis.

4.1 Data analysis and results

We analyzed the collected datasets by means of three measures: correlation, mutual information, and entropy (which is a particular instance of mutual information). For a more detailed description of the results see (Tarapore et al., 2004, 2005; Gómez et al., 2005).

4.2 Correlation

In the first behavioral state ("exploring"), the robot moved around avoiding obstacles and "searching" for red objects. The correlation matrix of a particular experimental run is displayed in Figure 5a.

In the second behavioral state ("tracking"), the robot moved toward the red object identified at the end of the previous state. In this case, the correlations between the activity of the red receptors in and close to the center of the image were high (see Fig. 6).

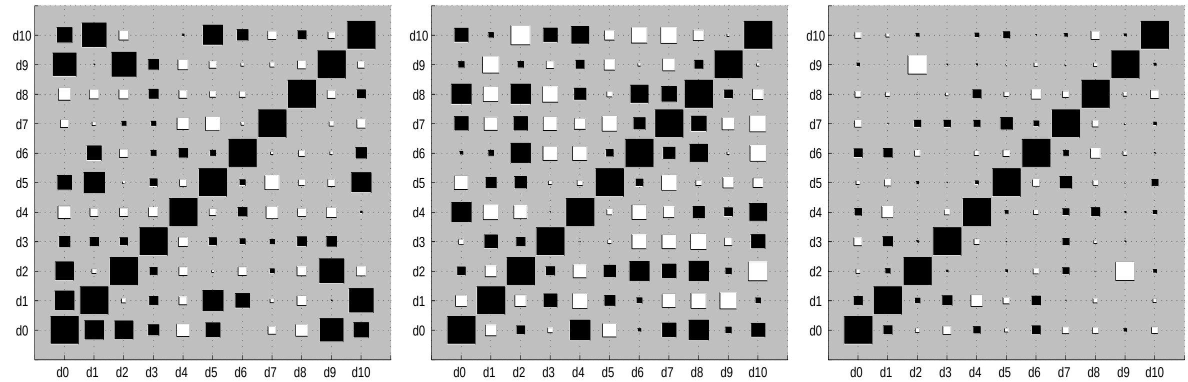

(a) (b) (c)

Figure 5: Correlation matrix obtained from the pair-wise correlation of the distance sensors for one particular experimental run. The behavioral states are: (a) "exploring", (b) "tracking", and (c) "circling." The higher the correlation, the larger the size of the square. White squares denote a negative correlation, black squares a positive one. The diagonal represents the auto-correlation of the individual time series. For the behavioral states displayed the average cross-correlation computed over 16 experimental runs was: 0.011 ± 0.004, 0.097 ± 0.012, and 0.083 ± 0.041, where $\pm$ indicates the standard deviation.

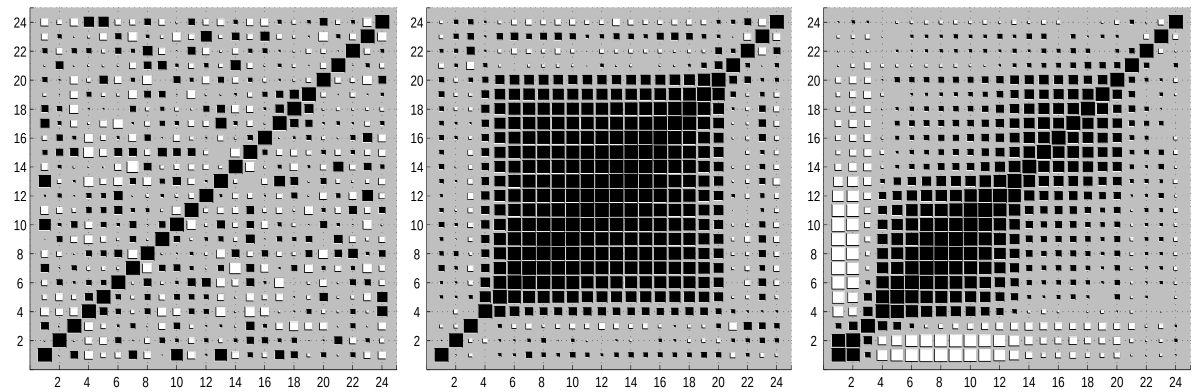

(a) (b) (c)

Figure 6: Correlation matrix obtained from the pair-wise correlation of the red channels for one particular experimental run. The behavioral states are: (a) "exploring", (b) "tracking", and (c) "circling." The higher the correlation, the larger the size of the square. White squares denote a negative correlation, black squares a positive one. The diagonal represents the auto-correlation of the individual time series. For the behavioral states displayed the average cross-correlation computed over 16 experimental runs was: 0.053 ± 0.023, 0.309 ± 0.042, and 0.166 ± 0.031, where $\pm$ indicates the standard deviation.

In the third behavioral state ("circling"), we observed negative correlations between the pairs of proximity sensors located on the ipsi-lateral (that is, same) side of the robot, such as (d_2, d_9) or (d_3, d_{10}).(see Fig. 6).

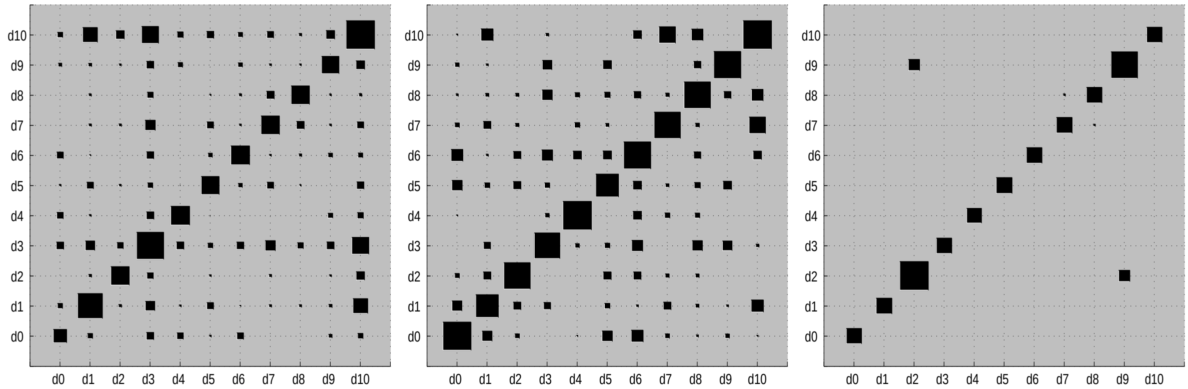

(a) (b) (c)

Figure 7: Mutual information matrix obtained by estimating the mutual information between pairs of proximity sensors in one particular experimental run during the behavioral state: (a) "exploring", (b) "tracking", (c) "circling". The higher the mutual information, the larger the size of the square. In order to better display the informational structure the plots have not been normalized. The maximum values for the mutual information are: 1.825 bits ("exploring"), 1.080 bits ("tracking"), and 2.936 bits ("circling").

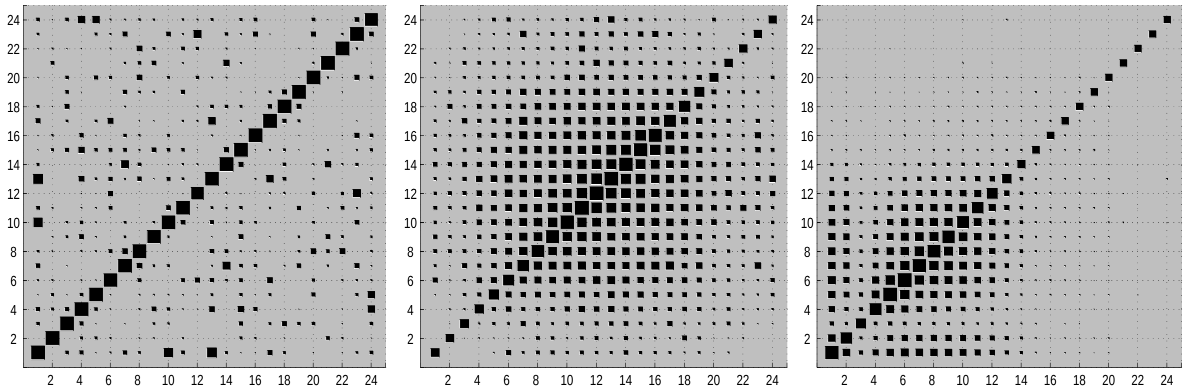

(a) (b) (c)

Figure 8: Mutual information matrix obtained by estimating the mutual information between pairs of red channels in one particular experimental run during the behavioral state: (a) "exploring", (b) "tracking", (c) "circling". The higher the mutual information, the larger the size of the square. In order to better display the informational structure the plots have not been normalized. The maximum values for the mutual information are: 0.999 bits ("exploring"), 3.128 bits ("tracking"), and 3.637 bits ("circling").

4.2.1 Entropy and mutual information

The pair-wise mutual information between the eleven proximity sensors is shown in Figure 7 and Figure 8 shows the mutual information matrices obtained from the estimation of the mutual information for pairs of red channels.

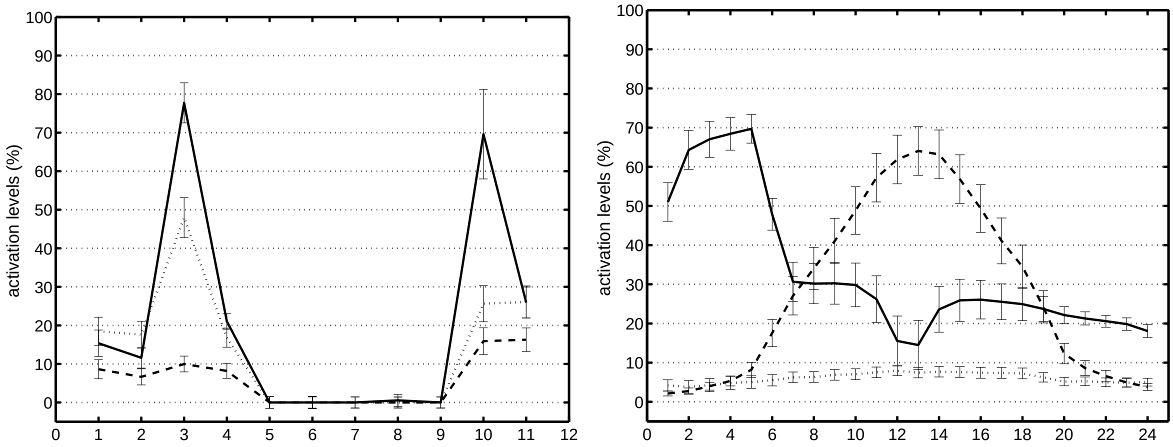

(a) (b)

Figure 9: (a) Plot of activation levels for the proximity sensors (d_0 to d_{10}) for the three behavioral states. (b) Plot of activation levels for the image sensors (i_1 to i_{24}) for the three behavioral states. The plots display the average computed over 16 experimental runs, the bars denoting the 95% confidence limit. The lines denote the following behavioral states: "exploring" (dotted), "tracking" (dashed), "circling" (continuous).

4.2.2 Cumulated sensor activation

The total sensory stimulation for both sensory modalities was computed by integrating – separately for each behavioral state – the activation of the individual sensors during an experimental run. The activation was then normalized as a percentage (see Figure 9).

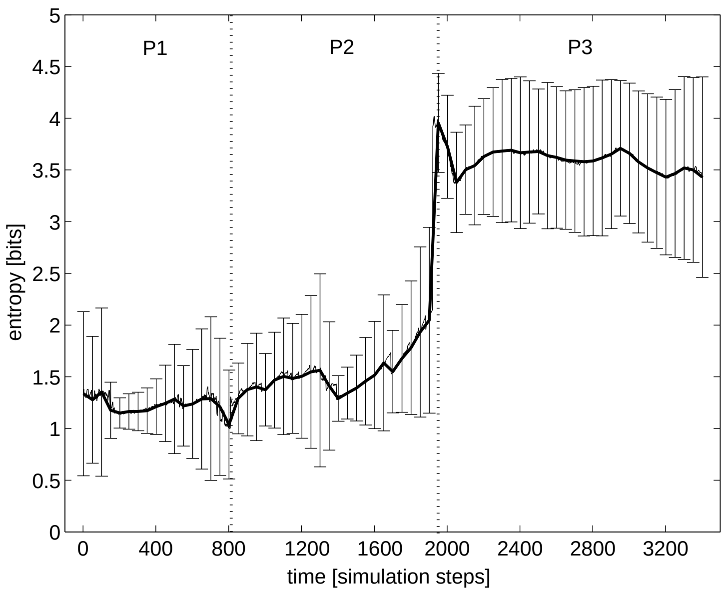

Figure 10: Entropy of the effective red color averaged over all vertical slices. The plots display the average computed over 16 experimental runs, the bars denoting the 95% confidence limit.

4.2.3 Pre-processed image entropy

The change over time of the total image entropy (computed as the average of the entropies of the individual vertical image slices i_{1-24}) is displayed in Fig. 10. While the robot explored its ecological niche (phase P_1), the image entropy was low and not highly variable (compared to P_2 and P_3). When the robot began approaching a red object (phase P_2), the image entropy started to increase. The image entropy reached its maximum in the third behavioral state, and stayed high as long as the robot kept circling around the red object (phase P_3).

5 Discussion

The information theoretic analyses of the sensory-motor data collected in several experiments performed with the 6 DOF robot arm and the active vision system as well as in simulation of a mobile robotic-agent show that (a) sensory-motor coordinated interaction with the environment can induce a significant reduction in the "complexity" of the input data, and can generate data that facilitate learning, and (b) that the particular temporal pattern of the correlations among the sensors and between these and the motor outputs can be used to characterize the robot-environment interaction. (Tarapore et al., 2004, 2005; Gómez et al., 2005). A freely available library of methods for quantifying the informational structure of sensory and motor data (written in MATLAB) can be downloaded from the following URL: http://sourceforge.net/projects/infometh/

References

Ballard, D. (1991). Animate vision. *Artificial Intelligence*, 48(1):57–86.

Cover, T. and Thomas, J. (1991). *Elements of Information Theory.* New York: John Wiley.

Gibson, E. (1988). Exploratory behavior in the development of perceiving, acting, and the acquiring of knowledge. *Annual Review of Psychology*, 39:1–41.

Gómez, G. and Eggenberger Hotz, P. (2004a). An evolved learning mechanism for teaching a robot to foveate. In *Proc. of the 9th Int. Symp. on Artificial Life and Robotics (AROB-9)*, pages 655–658.

Gómez, G. and Eggenberger Hotz, P. (2004b). Investigations on the robustness of an evolved learning mechanism for a robot arm. In *Proc. of the 8th Int. Conf. on Intelligent Autonomous Systems (IAS-8)*, pages 818–827.

Gomez, G., Hernandez, A., Eggenberger Hotz, P., and Pfeifer, R. (2005). (in press). an adaptive learning mechanism for teaching a robot to grasp. In *To appear in Proc. of AMAM 2005.*

Gómez, G., Lungarella, M., and Tarapore, D. (2005). Information-theoretic approach to embodied category learning. In *Proc. of the 10th Int. Symp. on Artificial Life and Robotics (AROB 10), Beppu, Oita, Japan.*, pages 332–337.

Lederman, S. J. and Klatzky, R. L. (1990). Haptic exploration and object representation. In Goodale, M., editor, *Vision and Action: The Control of Grasping*, pages 98–109. New Jersey: Ablex.

Lee, T. S. and Yu, S. X. (1999). An information-theoretic framework for understanding saccadic behaviors. In *Proc. of the First Intl. Conf. on Neural Information Processing*. Cambridge, MA: MIT Press.

Lungarella, M. and Pfeifer, R. (2001). Robots as cognitive tools: Information-theoretic analysis of sensory-motor data. In *In Proceedings of the 2nd International Conference on Humanoid Robotics, Waseda, Japan.*

Lungarella, M. and Sporns, O. (2004). Methods for quantifying the informational structure of sensory and motor data. *Neuroinformatics*. in preparation.

Nolfi, S. (2002). Power and limit of reactive agents. *Neurocomputing*, 49:119–145.

Pfeifer, R. and Scheier, C. (1997). Sensory-motor coordination: The metaphor and beyond. *Robotics and Autonomous Systems*.

Pfeifer, R. and Scheier, C. (1999). *Understanding Intelligence*. MIT Press.

Roulston, M. (1999). Estimating the errors on measured entropy and mutual information. *Physica D*, 125:285–294.

Scheier, C. and Pfeifer, R. (1997). Information theoretic implications of embodiment for neural network learning. In *ICANN 97*, pages 691–696.

Shannon, C. (1948). A mathematical theory of communication. *Bell System Tech. Journal*, 27.

Sporns, O. and Pegors, J. (2003). Generating structure in sensory data through coordinated motor activity. In *Proc. of Intl. Joint Conf. on Neural Networks*, page 2796.

Steuer, R., Kurths, J., Daub, C., Weise, J., and Selbig, J. (2002). The mutual information: detecting and evaluating dependencies between variables. *Bioinformatics*, 18:231–240. Suppl.2.

Tarapore, D., Lungarella, M., and Gómez, G. (2004). Fingerprinting agent-environment interaction via information theory. In *Proceedings of the 8th Int. Conference on Intelligent Autonomous Systems (IAS-8), Amsterdam, The Netherlands*, pages 512–520.

Tarapore, D., Lungarella, M., and Gomez, G. (2005). Quantifying patterns of agent-environment interaction. *Robotics and Autonomous Systems(accepted for publication).*

Te Boekhorst, R., Lungarella, M., and Pfeifer, R. (2003). Dimensionality reduction through sensory-motor coordination. In Kaynak, O., Alpaydin, E., Oja, E., and Xu, L., editors, *Proc. of the Joint Inlt. Conf. ICANN/ICONIP*, pages 496–503. LNCS 2714.

Thelen, E. and Smith, L. (1994). *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, MA: MIT Press. A Bradford Book.

Yarbus, A. (1967). *Eye movements and vision*. Plenum Press.