

Quantifying Patterns of Agent-Environment Interaction

Danesh Tarapore¹, Max Lungarella^{2,*} and Gabriel Gómez³

¹*School of Information Technology, Indian Institute of Technology, Bombay, India*

²*Lab. for Intelligent Systems and Informatics, The University of Tokyo, Japan*

³*Artificial Intelligence Lab., University of Zurich, Switzerland*

Abstract

This article explores the assumption that a deeper (quantitative) understanding of the information-theoretic implications of sensory-motor coordination can help endow robots not only with better sensory morphologies, but also with better exploration strategies. Specifically, 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 a simulated 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.

Key words: Self-structuring of information, sensory-motor coordination, agent-environment interaction

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

¹ Corresponding author: M.L. (maxl@isi.imi.i.u-tokyo.ac.jp)

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). Exploratory activity of hands and eyes is a particular instance of coordinated motor activity that extracts different kinds of information through interaction with the environment. In other words, robots and other agents are not passively exposed to sensory information, but they can actively shape such information. 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).

The experimental tool of the study presented in this article is a simulated robotic agent, which was programmed to search its local environment for red objects, approach them, and explore them for a while. The analysis of the recorded sensory and motor data shows that different types of sensory-motor activities displayed distinct fingerprints reproducible across many experimental runs. In the two following sections, we give a detailed overview of our experimental setup, and describe the actual experiments. Then, in Section 4,

we expose our methods of analysis. In Section 5, we present our results and discuss them. Eventually, in Section 6, we conclude and point to some future research directions.

2 Experimental Setup

The study was conducted in simulation. The experimental setup consisted of a two-wheeled robot and of a closed environment cluttered with randomly distributed, colored cylindrical objects. A bird’s eye view on the robot and its ecological niche is shown in Fig. 1 a. The robot was equipped with eleven proximity sensors (d_{0-10}) to measure the distance to the objects and a pan-controlled camera unit (image sensor) (see Fig. 1 b). The pan-angle of the camera was constrained to vary in an interval of $\pm 60^\circ$ relative to the agent’s midline. 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 (see caption of Fig. 1). 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. Negative values of R were set to zero. This operation guaranteed that the red channel gave maximum response for fully saturated red color, that is, for $r=31$, $g=b=0$. Subsequently, the red color slices will also be referred to as red channels or red receptors.

For the control of the robot, we opted for the Extended Braitenberg Architecture (Pfeifer and Scheier, 1999) (see Figure 1 c). In this architecture, each of the robot’s sensors is connected to a number of processes which run in parallel, continuously influencing the agent’s internal state, and governing its behavior. 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.” Two points are noteworthy. First, while the agent is exploring the environment, its behavioral activity is only “weakly” sensory-motor coordinated, in the sense that its motor activity is mainly driven by the random activity affecting its sensors. The two latter behavioral states, however, display “strong” sensory-motor coordinated activity, that is, motor activity characterized by a




Fig. 1. (a) Bird's eye view on the robot and its ecological niche. The trace depicts the path of the robot during a typical experiment. (b) Schematic representation of the simulated agent. The sensors have a position-dependent range: if rl is the length of the robot, 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 $0.6rl$. (c) Extended Braitenberg Control Architecture: As shown, four processes govern the agent's behavior.

tight coupling between sensing and acting. Second – and this point is more subtle – the segmentation of the observed behavior into distinct behavioral states is an important, maybe even necessary, step to simplify the analysis towards quantifying the agent-environment interaction and identifying stable patterns of interaction. It is crucial to understand that the goal here is to show how statistical and information-theoretic measures can be employed to identify patterns in sensory data, and not how those patterns can be identified automatically. In the latter case, despite being possible (see Fig. 7) the segmentation of the observed behavior in distinct states would make less sense.

3 Experiments

A top view of a typical experiment is shown in Fig. 1 a. At the outset of each experimental run, the robot’s initial position was set to the final position of the previous experiment (except 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 Methods

First, we introduce some notation. Correlation quantifies the amount of linear dependency between two random variables X and Y , and is given by $Corr(X, Y) = (\sum_{x \in X} \sum_{y \in Y} p(x, y) (x - m_X)(y - m_Y)) / \sigma_X \sigma_Y$, 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 (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 $H(X) = -\sum_{x \in X} p(x) \log p(x)$, 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 $H(X, Y) = -\sum_{x \in X} \sum_{y \in Y} p(x, y) \log p(x, y)$. 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 $MI(X, Y) = H(X) + H(Y) - H(X, Y)$. 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: $T = (B - 1)/2N$ 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 $T = (B_x + B_y - B_{x,y} - 1)/2N$ to the mutual information estimate (where B_x , B_y , $B_{x,y}$, and N have an analogous meaning to the previous case).

5 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). In this section we describe, and in part discuss, the results of our analyses.

5.1 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 2 a (the diagonal of the same plot represents the autocorrelation of d_i , that is, $Corr(d_i, d_i) = 1, i = 0..10$). In all performed experiments, we observed either no or only weak correlations between the proximity sensors, that is, the correlations were small and their absolute values close to zero. The cross-correlation averaged over 16 runs was 0.011 ± 0.004 , where 0.011 denotes the mean, and 0.004 is the standard deviation. Similarly, the output of the other sensory modality (red receptors) did not lead to stable correlation matrices, that is, the pair-wise correlations between the receptors varied significantly across the different experimental

runs (the correlation matrix of a particular run is shown in Figure 3 a). In this case, the average correlation was 0.053 ± 0.023 . A possible explanation for the absence of significant correlations in this state (in both sensory channels) is that the oscillatory movement of the robot’s camera induced a rapidly changing stream of sensory data devoid of correlative structure. The intrinsic white noise characterizing the sensor output in absence of external input, as well as the low predictability of the sensory activations while the robot was “randomly” exploring its ecological niche (i.e., the behavior was not sensory-motor coordinated), only exacerbated the difficulty in inducing statistical dependencies between the sensory channels.

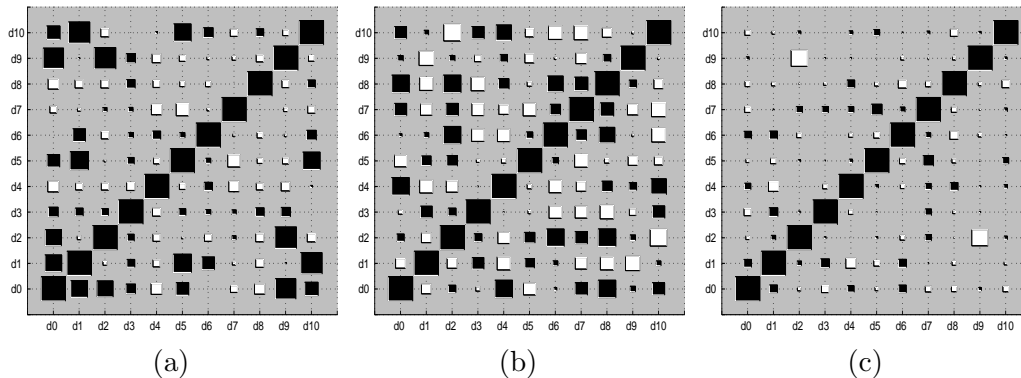


Fig. 2. Correlation matrix obtained from the pair-wise correlation of the distance sensors for one particular experimental run. The behavioral states are: (a) “exploring”, (b) “tracking”, and (c) “circling.” The higher the correlation, the larger the size of the square. White squares denote a negative 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.

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. 3 b), and the average correlation computed over 16 experimental runs, amounted to 0.309 ± 0.042 . A possible explanation is that the robot’s behavior in this “goal-corrected” state was largely governed by the dynamics of the sensory-motor coupling between robot and environment. In other words, the robot – while moving toward the object – kept correcting its movements so that the tracked object remained in the center of its visual field. Moreover, because this state was characterized by a goal-directed movement of the robot toward the red object, the number of red pixels present in the image increased, leading to an increase of the stimulation of the red receptors located in the center (note that the activation of the red receptors is an average computed over a vertical slice), and to a corresponding increase of the correlation between those receptors.




Fig. 3. 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}) . The correlation matrix shown in Figure 2c, for instance, displays such a negative correlation for the sensor pair (d_2, d_9) (the correlation is -0.671). The negative correlation finds explanation in the fact that the robot’s trajectory was not perfectly circular, but somewhat wobbly, thus leading to out-of-phase activation of the sensors in the front (d_9) and in the one in the back (d_2), and hence to a negative correlation. In this state, we also observed – in all performed experimental runs – strong correlations between the output of two groups of red receptors located off-center (see Figure 3c). In that particular experimental run, the correlation was 0.920 for the receptors $i_4 - i_{12}$ (a very high value indeed). The overall average (computed over 16 experimental runs) was 0.166 ± 0.031 . The evident asymmetry of the correlation matrix is a consequence of the limitations of the camera’s pan-angle ($\pm 60^\circ$), which caused the object to appear on the side and not in the center of the visual field.

5.2 Entropy and mutual information

The pair-wise mutual information between the eleven proximity sensors is shown in Figure 4. The diagonal of the same plot, courtesy of the expression $H(X) = MI(X, X)$, gives the entropy of the sensory stimulation. It is important to note that as a result of the additive uniform white noise affecting the activations of the individual sensors, even sensors that had never been active –

due to some object – were characterized by a large entropy (refer also to graph of cumulated activation in Fig. 6). A first observation is that in the first and




Fig. 4. 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”).

second behavioral states, the results of the analysis of the data gathered in a particular experiment cannot be generalized to all experiments (as opposed to correlation analysis). That is, the standard deviation between individual experimental runs was high. The reason is that in runs, in which the robot had to avoid some obstacle (i.e., a non-red object), the entropy of the individual sensors, as well as the average mutual information between adjacent sensors, were considerably larger than in experimental runs, in which the robot did not encounter any objects (it is important to remember that entropy measures the uncertainty of a sensory state). In the latter case, the entropy and mutual information were mainly due to the random basic activation of the sensors. Further, it is possible to observe that in the third behavioral state (“circling”), the entropy of the sensors d_2 and d_9 , as well as their mutual information were high, that is, $H(d_2) = 2.83$ bits, $H(d_9) = 2.75$ bits, and $MI(d_2, d_9) = 1.62$ bits (see Fig. 4 c). In contrast, the average mutual information was rather modest (0.243 ± 0.031 bits) (note that the highest possible value for entropy and mutual information is 5 bits).

Figure 5 shows the mutual information matrices obtained from the estimation of the mutual information for pairs of red channels. In the behavioral state (“exploring”), the average (entropy-devoided) mutual information computed over all experiments was 0.123 ± 0.020 bits (Fig. 5 a shows the result for one particular experiment). The reason for such low values of mutual information are (a) the additive white noise affecting the sensory channels, and (b) the oscillations of the camera from side to side, which led to a rapidly changing camera image, and consequently to a drop of the statistical dependence

between red channels (for a similar argument refer also to sub-section 5.1). In contrast to the mutual information, the entropy of the individual sensory channels was very high and close to the channel capacity, i.e., 5 bits. In the second behavioral state (“tracking”), the average entropy for the red receptors was not quite as high as in the first behavioral state, and amounted to (2.674 ± 0.362) bits). The same mutual information between the red receptors, however, was higher (0.804 ± 0.160) bits). The mutual information matrix for a particular experimental run is shown in Fig. 5 b). For qualitatively similar results refer also to (Lungarella and Pfeifer, 2001; Sporns and Pegors, 2003). In the third behavioral state (“circling”), the entropy of the peripheral red channels, as well as the mutual information between them, were large (see Fig. 5 c). Across all experiments, for both sides of the image sensor, the standard deviation of the mutual information assumed high values (e.g., in Fig. 5 c, the standard deviation of the six receptors on the far left of the image sensor is 0.561 bits; the robot was circling around the object counter-clockwise). In contrast, the mutual information for the red channels close to the center was low (0.244 bits). This low value was largely independent from the direction in which the robot circled around the object. The mutual information between red receptors across all the experiments was also low: 0.102 ± 0.020 bits.

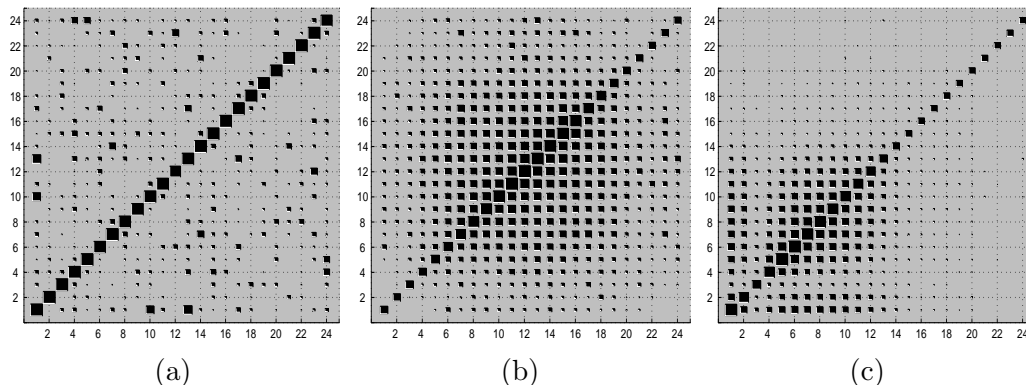


Fig. 5. Mutual information matrix obtained by estimating the mutual information between pairs of red channels in one particular experimental run during the behavioral state: (a) “exploring”, (b) “tracking”, (c) “circling”. The higher the mutual information, the larger the size of the square. In order 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”).

5.3 Cumulated sensor activation

It is important not to confuse the amount of variability (information) with the cumulated amount of activation (total stimulation) of a particular sensor. The total sensory stimulation for both sensory modalities was computed by inte-

grating – separately for each behavioral state – the activation of the individual sensors during an experimental run. The activation was then normalized as a percentage (see Fig. 6). Interestingly, in the behavioral states “exploring” and “tracking” the cumulated sensor activation did not show any stable patterns across multiple experiments, in the sense that the positions of the peaks changed from experiment to experiment, and were dependent on the number of objects encountered. In the third behavioral state (“circling”), however, the activation levels of the sensors d_2 and d_3 were high and stable across all experimental runs (depending on the direction the agent circled around the object). Figure 6(a) demonstrates such stability by displaying the average (and the standard deviation, that is, variance) of a set of experimental runs in which the robot circled around the object counter-clockwise. The same graph shows that the sensor d_{10} was also characterized by large activation levels (again, for the case in which the agent circled counter-clockwise around objects).

We computed the activation levels also for the 24 red receptors. As shown in Figure 6 b, the total stimulation of the red channels in the first behavioral state (“exploring”) was low in all experiments (the stimulation was largely due to noise in the sensory channels). As in the case of the distance sensors, the graphs are averages over a set of experimental runs in which the robot circled counter-clockwise around objects. In the second behavioral state (“tracking”), however, the activation levels for the red receptors close to the center were high, and gradually tapered out toward the periphery. Such decrease of activation is a result of the continuous sensory-motor adjustments of the camera pan-angle to keep the red object in the center of its visual field. Clearly, the peripheral red receptors are barely stimulated. Finally, the third behavioral state (“circling”) shows high activation levels for the image sensors on both sides of the robot (for the data displayed in Figure 6 b the agent circled counter-clockwise around the object).

5.4 *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. 7. 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).

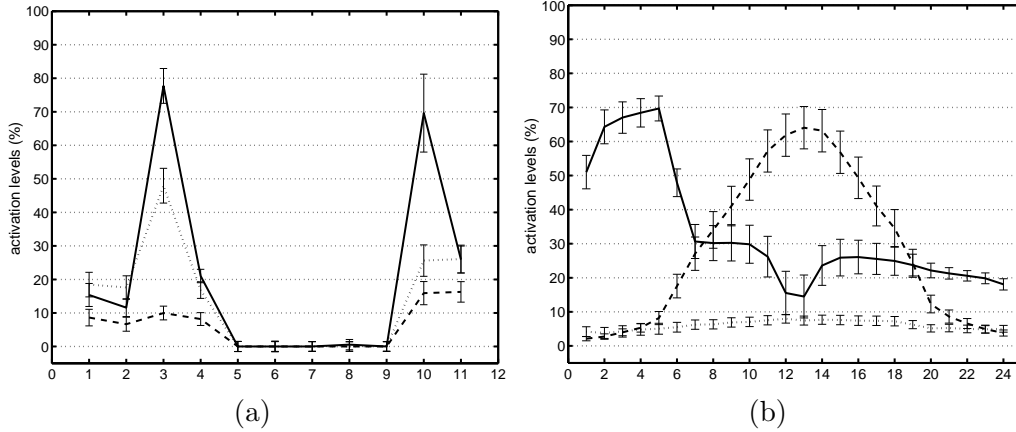


Fig. 6. (a) Plot of activation levels for the proximity sensors (d_0 to d_{10}) for the three behavioral states. (b) Plot of activation levels for the image sensors (i_1 to i_{24}) for the three behavioral states. The plots display the average computed over 16 experimental runs, the bars denoting the 95% confidence limit. The lines denote the following behavioral states: “exploring” (dotted), “tracking” (dashed), “circling” (continuous).




Fig. 7. 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.

6 Further Discussion and Conclusion

To summarize, coordinated motor activity leads to correlations in the sensory data that can be used to characterize the robot-environment interaction. Statistical measures, such as correlation and mutual information (although sometimes tricky to use), can be effectively employed to extract and quantify patterns induced by the coupling of robot and environment. In the “circling” behavioral state, for instance, the average correlation (evaluated over 16 exper-

imental runs) normalized by the number of distance sensors or red receptors was 0.083 ± 0.041 for the distance sensors and 0.166 ± 0.031 for the image sensors. The “tracking” behavioral state displayed an even higher stability, and the average correlation was 0.097 ± 0.012 (for the distance sensors) and 0.309 ± 0.042 (for the image sensors). Mean and standard deviation clearly show that correlation analysis leads to stable sensory patterns across multiple experimental runs, and thus define quantitative measures for behavioral fingerprints. A similar results holds for mutual information. It is important to note that such stability is the direct consequence of a sensory-motor coordinated interaction. And indeed, for the case in which the agent tracked the object (in this case the agent’s sensory-motor coupling with the environment was strong) the fingerprints displayed least variance. A possible conclusion is that the particular combination of (a) morphological setup (that is, two wheels, 24 image sensors, and 11 distance sensors) and (b) control architecture is well suited for tracking objects.

Our analyses also show that although correlation and mutual information provide appropriate statistical measures for fingerprinting interaction, they differ in at least one important aspect. Correlation can be used to identify fingerprints of robot-environment interaction only if the sensory activations between different sensors are temporally contiguous (that is, temporally close relative to the time scale of the agent’s control structure). We hypothesize that such temporal contiguity in the raw (that is, unprocessed) sensory data can be induced by coordinated motor activity. It may indeed be the case that sensory-motor coordination provides a very natural mechanism to achieve a matching of the various time scales affecting an agent’s behavior: environment, neural, and body-related.

A further result (that will be elaborated in future work) is that even if the sensory channels are affected by additive white noise, adequate sensory-motor coordination interaction can indeed induce stable fingerprints. In this sense, it is possible to put forward the hypothesis that sensory-motor coordinated interaction can act as some sort of “behavioral denoising filter.”

7 Acknowledgments

Max Lungarella would like to thank the University of Tokyo, and a Special Coordination Fund for Promoting Science and Technology from the Ministry of Education, Culture, Sports, Science, and Technology of the Japanese government. For Gabriel Gómez funding has been provided by grant number NF-10-101827/1 of the Swiss National Science Foundation and the EU-Project ADAPT (IST-2001-37173).

References

- Ballard, D., 1991. Animate vision. *Artificial Intelligence* 48 (1), 57–86.
- Cover, T., 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.
- Lederman, S. J., Klatzky, R. L., 1990. Haptic exploration and object representation. In: Goodale, M. (Ed.), *Vision and Action: The Control of Grasping*. New Jersey: Ablex, pp. 98–109.
- Lee, T. S., 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., Pfeifer, R., 2001. Robots as cognitive tools: Information-theoretic analysis of sensory-motor data. In: *Proc. of the 2nd Int. IEEE/RSJ Conf. on Humanoid Robotics*. pp. 245–252.
- Lungarella, M., 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., Scheier, C., 1997. Sensory-motor coordination: The metaphor and beyond. *Robotics and Autonomous Systems* 20, 157–178.
- Pfeifer, R., Scheier, C., 1999. *Understanding Intelligence*. Cambridge, MA: MIT Press.
- Roulston, M., 1999. Estimating the errors on measured entropy and mutual information. *Physica D* 125, 285–294.
- Scheier, C., Pfeifer, R., 1997. Information theoretic implications of embodiment for neural network learning. In: *ICANN 97*. pp. 691–696.
- Shannon, C., 1948. A mathematical theory of communication. *Bell System Tech. Journal* 27.
- Sporns, O., Pegors, J., 2003. Generating structure in sensory data through coordinated motor activity. In: *Proc. of Intl. Joint Conf. on Neural Networks*. p. 2796.
- Steuer, R., Kurths, J., Daub, C., Weise, J., Selbig, J., 2002. The mutual information: detecting and evaluating dependencies between variables. *Bioinformatics* 18, 231–240, suppl.2.
- Te Boekhorst, R., Lungarella, M., Pfeifer, R., 2003. Dimensionality reduction through sensory-motor coordination. In: Kaynak, O., Alpaydin, E., Oja, E., Xu, L. (Eds.), *Proc. of the Joint Intl. Conf. ICANN/ICONIP*. pp. 496–503, INCS 2714.
- Thelen, E., 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.