Fingerprinting Agent-Environment Interaction Via Information Theory

Danesh Tarapore[1], Max Lungarella[2,*] and Gabriel Gómez[3]

[1]*School of Information Technology, Indian Institute of Technology, Bombay, India*
[2]*Neuroscience Research Institute, Tsukuba AIST Central 2, Japan*
[3]*Artificial Intelligence Laboratory, University of Zurich, Switzerland*

Abstract. In this paper, 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. We 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 quantify (fingerprint) the agent-environment interaction, and (c) to analyze the informational relationship between the different components of the sensory-motor apparatus. We hypothesize that a deeper understanding of the information-theoretic implications of sensory-motor coordination can help us endow our robots with better sensory morphologies, and with better strategies for exploring their surrounding environment.

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 [3]. 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 to areas of a visual scene or an image that deliver useful and essential perceptual information [14]. To reason about the organization of saccadic eye movements, Lee and Yu [4] 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.

Underlying these perceptual abilities there is a process of sensory-motor coordination that couples action and perception. Thus, coordinated movements must be considered part of the perceptual system [13], and whether the sensory stimulation is visual, tactile, or auditory, perception always includes associated movements of eyes, hands, arms, head and neck [1, 2].

[2,*] Corresponding author: max.lungarella@aist.go.jp

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 [5, 10]. 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 that 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) characterizing the agent-environment interaction. Our approach builds on top of previous work on category learning [6, 8], as well as on information-theoretic and statistical analysis of sensory-motor data [5, 10, 12].

In this paper, we simulated a robotic agent whose task was to search its surrounding environment for red objects, approach them, and explore them for a while. The analysis of the recorded sensory-motor data showed that different types of sensory-motor activities displayed distinct fingerprints reproducible across many experimental runs. In the two following sections, we give an overview of our experimental setup, and describe the actual experiments. In Section 5, we present our results and discuss them. Eventually, we conclude and point to some future research directions.

2 Experimental Setup

We conducted our study 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 (left). The robot was equipped with 11 proximity (distance) sensors (d_{0-10}) and a pan-controlled camera unit (image sensor) – see Fig.1 (center). The proximity sensors had a position-dependent range, that is, the sensors in the front and 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, and was partitioned into a space having 32 discrete states ($5\,bit$ sensor resolution). To reduce the dimensionality of the input data, we divided the camera image into 24 vertical rectangular slices with widths decreasing toward the center. Then we computed the amount of "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 a maximum in the response of the red channel for fully saturated red color, i.e., for *r=31, g=b=0*. 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 [7]. In this architecture, each of the robot's sensors is connected to a number of processes which run in parallel and continuously influence the agent's internal state, as well as 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." It is important to note that three behavioral states display coordinated motor activity, and are characterized by a tight coupling between sensing and acting. We advance that the

Figure 1: (Left) Bird's eye view on the robot and its ecological niche. The trace represents a typical path of the robot during one experiment. (Center) Schematic representation of the agent. The sensors have a position-dependent range. Being rl the length of the robot, the range of d_0, d_1, d_9, and d_{10} is $1.8\,rl$, the one of d_2 and d_3 is $1.2\,rl$, and the one of d_4, d_5, d_6, d_7, and d_8 is $0.6\,rl$. (Right) Extended Braitenberg Control Architecture.

segmentation of the observed behavior into distinct behavioral states is an important (maybe even necessary) step for fingerprinting the agent-environment interaction, i.e., for identifying stable patterns of interaction (such as stereotyped exploratory hand movements).

3 Methods

First, we introduce some notation. Correlation quantifies the amount of linear dependency between two random variables X and Y, 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 $p(x)$. Analogously, 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 for mutual information, we assumed the binary logarithm. Both are measure in bit. The entropy of the sensory data lies between the minimum ($0\,bit$) and maximum entropy bound ($5\,bit$). 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 [11]. 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 [11]. Because the sensors had a resolution of $5\,bit$, we estimated the histogram by setting the number of bins to 32 (equivalent 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 estimation of $p(x, y)$, thus speeding up the computation. As noted previously, the distance sensors are identified by d_i, $i \in [0, 10]$, whereas the effective red color sensors are indexed with the numbers 1 to 24.

4 Experiments

At the outset of each experimental run, the robot's initial position was set to the end position of the previous experiment 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 camera stopped rotating from side to side, and the robot started moving in the direction identified by the camera orientation, 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 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. A top view of a typical experiment is shown in Fig.1 (left). We conducted 16 experiments. Each experiment consisted of approximately 1400 data samples, which were stored into a time series file for subsequent analysis.

5 Data Analysis and Results

We analyzed the collected datasets by means of three measures: correlation, mutual information, and entropy (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 moves around avoiding obstacles and "searching" for red objects. In all performed experiments, there were either no or only weak correlations between the proximity sensors, that is, the correlations were small and their absolute value close to zero. In Fig.2 (left), for instance, the average correlation is 0.011. The intrinsic noise of the sensors, as well as the unpredictability of the sensory activations while the robot is exploring its ecological niche, make the identification of statistical dependencies between the sensory activations by means of linear correlation difficult. Similarly, the output of the red channels do not lead to a "stable" correlation matrix, that is, the pair-wise correlations between the sensory channels varies significantly between the different experimental runs. The average correlation in the case of Fig.3 (left) is 0.053 (again a low value), and the standard deviation is 0.023. The reason is that in this state, the oscillatory movement of the robot's camera induces a rapidly changing stream of sensory data, and consequently leads to small correlations between the red channels.

In the second behavioral state ("tracking"), the robot moves toward the previously identified red object (see Fig.2 center). In this case, the correlations between the activity of the red receptors in and close to the center of the image are high (see Fig.3 center). A possible explanation is that the robot keeps correcting the direction of its movements so that the tracked object remains in the center of its visual field. Moreover, because this state is characterized by

a goal-directed movement of the robot toward the red object, the number of red pixels present in the image increases, leading to an increase of the stimulation of the red receptors located in the center (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.

In the third behavioral state ("circling"), we observed negative correlations (-0.442) between the pairs of proximity sensors located on the ipsi-lateral (same) side of the robot, such as (d_2, d_9) or (d_3, d_{10}) (see Fig.2 right). Due the non-linearities of the data and the noise-induced correlations, however, these correlations are not immediately evident from the plot. In this state, we observed in all performed experimental runs, strong correlations between the output of the red channels located in (and close to) the central image area (see Fig.3 right). The correlation was 0.92 for receptors in the center, with an overall average of 0.1658. The standard deviation of the correlation computed over all experiments was 0.0412. While circling around the object, the robot kept foveating on it. Due to the limitations of the camera angle, however, the object appears on the side and not in the center of the field of view.


Figure 2: Correlation matrix obtained from the pair-wise correlation of the distance sensors for one particular experimental run. From left to right the behavioral states are: "exploring", "tracking" and "circling." The higher the correlation, the larger the size of the square.


Figure 3: Correlation matrix obtained from the pair-wise correlation of the red channels for one particular experiment. From left to right the behavioral states are: "exploring", "tracking" and "circling." The higher the correlation, the larger the size of the square.

5.2 Entropy and mutual information

The pair-wise mutual information between the 11 proximity sensors is shown in Fig.4. The diagonal of the same plot gives the entropy of the sensory stimulation (remember the expression $H(X) = MI(X, X)$). Because the individual sensors are affected by uniform white noise, even sensors that are never active (see graph of cumulated activation during an experimental run Fig.6), can be characterized by a potentially large entropy.

In the first and second behavioral states, the results of the analysis for the data gathered in a particular experiment cannot be generalized to all experiments. The reason is that in experiments in which the robot avoids obstacles, the average mutual information between sensors, as well as the entropy of the individual sensors, is larger compared to experimental runs in which the robot does not encounter any object. In the third behavioral state "circling", the entropy of the activation of the sensors on both sides of the robot is large: $H(d_3) = 2.83\,bit$ and $H(d_{10}) = 2.75\,bit$ (see Fig.4 (right)). In the same figure, the mutual information between these sensors is also high: $MI(d_2, d_9) = 0.62\,bit$.

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 mutual information computed over all experiments is $0.123\,bit$, and the standard deviation is $0.020\,bit$ (Fig.5 (left) shows the result for one particular experiment). The reason for the low values of mutual information is that the camera oscillates from side to side, thus leading to a rapidly changing camera image, and hence to a drop of the statistical dependence between red channels. In the second behavioral state "tracking", the entropy for the red receptors in and around the center is high in comparison with the one of the first behavioral state (mean: $2.674\,bit$, standard deviation: $0.362\,bit$). The same holds for the mutual information between the red receptors (mean: $0.604\,bit$, standard deviation: $0.160\,bit$) (see Fig.5 center). In the third behavioral state, the entropy of the red channels at the periphery, as well as the mutual information between them, is large (see Fig.5 right). Across all experiments, for both sides of the image sensor, the standard deviation of the mutual information assumes high values (e.g., the standard deviation of the receptor on the far left of the image sensor is $0.461\,bit$). In contrast, the standard deviation for the red channels close to the center is low (e.g., $0.244\,bit$), and largely independent from the direction in which the robot is moving around the object. The standard deviation in the mutual information between red receptors across all the experiments was low ($0.102\,bit$). We conclude that mutual information may provide a good and stable measure for identifying and characterizing agent-environment interaction.

5.3 Cumulated sensor activation

The amount of variability (information) should not be confused with the cumulated amount of sensory activation (total stimulation) of a particular sensor. 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 experiment. We then normalized the activation as a percentage (see Fig.6). In the "exploring" and "tracking" behavioral states the cumulated sensor activation does not show any stable patterns across multiple experiments, in the sense that the positions of the peaks change from experiment to experiment and depend on the number of objects encountered. In the third behavioral state, however, the activation levels of the sensors d_2 and d_3 are high and stable across all experimental runs (see Fig.6 left). These sensors are used when the robot moves toward the red object. The same graph

Figure 4: Mutual information matrix obtained by estimating the mutual information between pairs of proximity sensors in one particular experimental run. From left to right the behavioral states are: "exploring", "tracking", and "circling". The higher the mutual information, the larger the size of the square.


Figure 5: Mutual information matrix obtained by estimating the mutual information between pairs of red channels in one particular experimental run. From left to right the behavioral states are: "exploring", "tracking", and "circling". The higher the mutual information, the larger the size of the square.

shows that the activation levels of the sensors d_9 or d_{10} are characterized by large values. These particular sensors are used to prevent the robot from colliding with the object (while circling around it). As for the distance sensors, we also computed the activation levels of the 24 red receptors (see Fig.6 center). The total stimulation of the red channels in the first behavioral state displays no stability across all experiments. In the second behavioral state the activation levels for the red receptors close to the center are high. The activation levels, however, gradually decrease toward the periphery. The decrease is a result of the continuous adjustments of the camera pan-angle in order to keep the red object in the center of its visual field. Thus, the peripheral red receptors are not stimulated. The behavioral state "circling" shows high activation levels for the image sensors on both sides of the robot.

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 slices) is displayed in Fig.6 (right). While the robot is exploring its ecological niche, the image entropy is low and constant (phase P_1), that is, there is not much variability in the sensory channel. When the robot starts approaching the red object (second behavioral state), the image entropy begins to increase (phase P_2). The image entropy reaches

its maximum in the third behavioral state, and stays high as long as the robot keeps circling around the red object (phase P_3).


Figure 6: (Left) Plot of activation levels for the proximity sensors for the three behavioral states. (Center) Plot of activation levels for the image sensors (1 to 24) for the three behavioral states. (Right) Entropy of the effective red color averaged over all vertical slices. Dashed line or P_1: exploring; dotted line or P_2: tracking; solid line or P_3: circling.

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, can be employed to extract fingerprints of the robot-environment interaction. In the "circling" behavioral state, for instance, the average correlation (evaluated over 16 experimental runs) divided by the number of distance sensors (11) or red receptors (24) is 0.083 ± 0.0412 for the sensors and 0.1658 ± 0.031 for the receptors (where $\pm$ indicates the standard deviation). Mean and standard deviation clearly show that the fingerprint (extracted by means of correlation analysis) is stable across multiple experimental runs. Similarly, in the "tracking" behavioral state, the average correlation is 0.097 ± 0.012 (for the distance sensors) and 0.309 ± 0.042 (for the red receptors). These results hold also for the mutual information.

Although correlation and mutual information both provide appropriate statistical measures for fingerprinting interaction, they differ in at least one important aspect. Whereas correlation can be used to identify fingerprints of robot-environment interaction only if the sensory activations between different sensors happen to be temporally contiguous, mutual information reveals nonlinear dependencies between the sensory stimulations that correlation cannot capture. This difference seems to be less of a problem if the interaction between the agent and its local environment is sensory-motor coordinated. Hence our hypothesis that temporal contiguity and stability in the raw sensory data are the result of coordinated motor activity (exploration strategy). In other words, sensory motor coordination actively generates input data containing high amounts of informational structure. Our analyses demonstrate that even if the sensory channels are affected by additive white noise, a proper sensory-motor coordinated interaction can indeed lead to stable fingerprints. High entropy could be the consequence of a "complex" robot behavior. The high entropy values correspond to more uncertainty and therefore more interesting behaviors. This would explain the high entropy values of the image sensors in the behavioral states "tracking" and "circling" as compared to the behavioral state "exploring". The mutual information gives the amount of information

shared by different sensors. Sensors coordinating with the motor in a particular behavioral state exhibit high mutual information. We conclude that the information shared by sensors and motors provides a fingerprint for the corresponding behavior.

Acknowledgments

References

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

[2] Gibson, E.J. Exploratory behavior in the development of perceiving, acting, and the acquiring of knowledge. *Annual Review of Psychology*, 39, p.1-41, 1988.

[3] Lederman, S.J. and Klatzky, R.L. Haptic exploration and object representation. In M. Goodale (ed.), *Vision and Action: The Control of Grasping*, p.98-109, NJ: Ablex, 1990.

[4] Lee, T.S. and Yu, S.X., An information-theoretic framework for understanding saccadic behaviors. In *Proc. of the 1st Int. Conf. on Neural Information Processing*, 1999.

[5] Lungarella, M. and Pfeifer, R. Robots as cognitive tools: Information-theoretic analysis of sensory-motor data. In *Proc. of the 2nd Int. IEEE/RSJ Conf. on Humanoid Robotics*, p.245-252, 2001.

[6] Pfeifer, R and Scheier, C. Sensory-motor coordination: The metaphor and beyond. *Robotics and Autonomous Systems*, 20, p.157-178, 1997.

[7] Pfeifer, R. and Scheier, C. *Understanding Intelligence.* Cambridge, MA: MIT Press, 1999.

[8] Scheier, C. and Pfeifer, R. Information-theoretic implications of embodiment for neural network learning. In *Proc. of Int. Conf. on Artificial Neural Networks*, p.691-696, 1997.

[9] Shannon, C. A mathematical theory of communication. In *Bell System Tech. Journal*, 27, 1948.

[10] Sporns, O. and Pegors, T. Generating structure in sensory data through coordinated motor activity. In *Proc. of Int. Joint Conf. on Neural Networks*, p.2796, 2003.

[11] Steuer, R., Kurths, J., Daub, C.O., Weise, J. and Selbig, J. The mutual information: detecting and evaluating dependencies between variables. *Bioinformatics*, 18, Suppl.2, p.231-240, 2002.

[12] Te Boekhorst, R., Lungarella, M. and Pfeifer, R. Dimensionality reduction through sensory-motor coordination. In *Proc. of the Joint Int. Conf. ICANN/ICONIP*, LNCS 2714, p.496-503, 2003.

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

[14] Yarbus, A.L. *Eye Movements and Vision.* Plenum Press, 1967.