

Sensor Adaptation and Development in Robots by Entropy Maximization of Sensory Data*

Lars Olsson*, Chrystopher L. Nehaniv*[†], Daniel Polani*[†]

Adaptive Systems and Algorithms[†] Research Groups*

School of Computer Science

University of Hertfordshire

College Lane, Hatfield Herts AL10 9AB

United Kingdom

{L.A.Olsson, C.L.Nehaniv, D.Polani}@herts.ac.uk

Abstract—A method is presented for adapting the sensors of a robot to the statistical structure of its current environment. This enables the robot to compress incoming sensory information and to find informational relationships between sensors. The method is applied to creating sensoritopic maps of the informational relationships of the sensors of a developing robot, where the informational distance between sensors is computed using information theory and adaptive binning. The adaptive binning method constantly estimates the probability distribution of the latest inputs to maximize the entropy in each individual sensor, while conserving the correlations between different sensors. Results from simulations and robotic experiments with visual sensors show how adaptive binning of the sensory data helps the system to discover structure not found by ordinary binning. This enables the developing perceptual system of the robot to be more adapted to the particular embodiment of the robot and the environment.

Index Terms—Ontogenetic robotics, sensory systems, entropy maximization

I. INTRODUCTION

One of the major tasks of many sensory processing system is compression of incoming sensory signals to representations more suitable to compute the specific quantities needed for that specific animal or robot to function in the world. It is believed that in many animals the functionality of the sensory organs and nervous system is almost completely innate, while in others it develops during the lifetime of the individual [4]. This development and adaptation is in part dependent on the structure of the incoming sensory signals, and there are also indications that individual neurons adapt to the statistical structure of their incoming sensory signals. This paper presents a robotic system that constantly adapts its visual sensors to the statistical structure of its environment by entropy maximization of the incoming sensory data.

The structure of the incoming sensory signals depends on the embodiment and actions of the agent and the environment. Research into the structure of natural signals is still at an early and explorative phase, but there are indications

that signals of different sensory modalities such as acoustic waveforms, odor concentrations, and visual contrast share some statistical properties [17]. For example, in general the local contrast in natural images has the same exponential form of the probability distributions as sound pressure in musical pieces [17]. Another commonality between signals of different modalities is coherence over time and, in many cases, spatial coherence. Coherence between signals means that one part of the signal can be predicted by another part of the signal. In other words, natural signals contain some *redundancy*. For example, consider nearby photoreceptors, which usually sample regions of visual space close to each other. Thus, nearby photoreceptors often sample the same object in natural scenes which usually is coherent in respect to colour, orientation, and other parameters. Contrast this with an image where each pixel is generated independently from a random distribution. An image like this will contain no redundancy.

Given the statistical structure and redundancy of natural signals is it natural to ponder whether this structure is exploited by animals to optimize their sensory systems. Barlow suggested in 1961 [1] that the visual system of animals “knows” about the structure of natural signals and uses this knowledge to represent visual signals. Thus, the sensory data can be represented in a more efficient way than if no structure in the data is known. In 1981 Laughlin recorded the distribution of contrasts as seen through a lens with the aperture the size of a fly photoreceptor while moving in a forest [8]. A single cell in the fly encodes contrast variations with a graded voltage response. The distribution of contrasts has some shape and Laughlin was interested in whether the voltage response conveyed the maximal amount of information given the specific distribution of contrasts by maximizing the entropy of the voltage distribution. This can be viewed as single neuron cell version of Linsker’s Infomax principle [9]. Laughlin compared the computed ideal conversion of contrast to voltage given his collected data from the forrest and found the match to be very good with the measured response of the second order neurons in the fly visual system. This result suggests that the early visual system of the fly is adapted to the statistical structure of natural scenes in its environment. Since Laughlin’s work focused on

*The work described in this paper was partially conducted within the EU Integrated Project RobotCub (“Robotic Open-architecture Technology for Cognition, Understanding, and Behaviours”) and was funded by the European Commission through the E5 Unit (Cognition) of FP6-IST under Contract FP6-004370.

global statistics the adaptation must have taken place over evolutionary time. Recent results indicate that the fly visual system also adapts to the current conditions in much shorter timescales, on the order of seconds or minutes [3]. This means that individual neurons adapt their input/out relations depending the structure of incoming signals.

In many robotic systems the processing of sensory data is often, like in the visual system of organisms, limited by the physical limits of sensation [2], memory capacity, processing speed, heat generation, power consumption [6], and limited bandwidth of data transfer. Thus, there is a need for robots to extract *relevant* information [11], [16] from the incoming streams of sensory data and also to represent this information as efficiently as possible. One way information can be represented as efficiently as possible given memory and processing constraints is by maximization of the entropy in each sensory channel as described above in the case of the fly. This is a well known statistical technique also known as *adaptive binning* [20] or *histogram* or *sampling equalization* [10]. This maximization also enables the robot to find more structure in ensembles of sensors since the maximization conserves the correlations between sensors while more resolution is achieved in the parts of the input distribution where most of the data is located.

This paper describes a sensory system that maximizes the information a robot with limited computational resources can have about the world. The sensory system is constantly adapting to the structure of its current environment using entropy maximization of each sensor using a sliding window mechanism. We show in simulation how an agent using this method can find informational relationships in the sensory data using the sensory reconstruction method [13] not found by a non-adapting system using the twice the amount of memory to represent the data. We also present results from experiments using a SONY AIBO¹ robot. The results show how the visual signals in different natural environments have different statistics and how the adaptive binning method helps the developing robot to reconstruct its visual field.

The structure of the rest of this paper is as follows. The next section describes the idea of entropy maximization and the information theory background. In section III the sensory reconstruction method is described and section IV presents the performed simulations and robotic experiments. Finally, section V concludes and points out some possible future areas of research.

II. ENTROPY MAXIMIZATION OF SENSORY DATA

To get a better understanding of entropy maximization, this section contains a short introduction to the general concepts of entropy and information theory [18]. Then entropy maximization is introduced and exemplified.

A. Information Theory

Let \mathcal{X} be the alphabet of values of a discrete random variable (information source, in this paper a sensor) X with

a probability mass function $p(x)$, where $x \in \mathcal{X}$. Then the *entropy*, or uncertainty associated with X is

$$H(X) = - \sum_{x \in \mathcal{X}} p(x) \log_2 p(x) \quad (1)$$

and the *conditional entropy*

$$H(Y|X) = - \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} p(x, y) \log_2 p(y|x) \quad (2)$$

is the uncertainty associated with the discrete random variable Y if we know the value of X . In other words, how much more information do we need to fully predict Y once we know X .

The *mutual information* is the information shared between the two random variables X and Y and is defined as

$$I(X; Y) = H(X) - H(X|Y) = H(Y) - H(Y|X). \quad (3)$$

To measure the dissimilarity in the information in two sources Crutchfield's information distance [5] can be used. The *information metric* is the sum of two conditional entropies, or formally

$$d(X, Y) = H(X|Y) + H(Y|X). \quad (4)$$

Note that X and Y in our system are information sources whose $H(Y|X)$ and $H(X|Y)$ are estimated from the time series of two sensors using (2).

B. Entropy Maximization

Due to memory and processing constraints, as well as to simplify learning, it is often desirable to compress incoming sensory data. One common method to achieve this is binning, whereby the range of incoming data is mapped to a smaller number of values using a transfer function. For example, consider the grey-scale image depicting Claude Shannon in fig. 1(a) where each pixel can have a value between 0 and 255. How could this image be compressed if only 5 different pixel values were allowed? Maybe the first method that comes to mind is to divide the range $\{0, 1, \dots, 255\}$ into 5 bins of size 51, where all values between 0 and 50 would be encoded as 0, 51 to 102 as 1, and so forth. This method, which does not take into account the statistics of the data, is called uniform binning, and the corresponding image is shown in fig. 1(c). As seen in fig. 1(d) the distribution of grey-scales in fig. 1(a) is not uniform, with most pixels in the range $\{100, 101, \dots, 200\}$. The entropy of the encoded image is ≈ 1.97 , which is less than the maximal theoretical entropy of $\log_2 5 \approx 2.32$. From an information theoretical point of view this means that this encoding is non-optimal since the entropy of the encoded image is less than the maximal possible entropy of the image. Now, consider fig. 1(e) which also uses 5 bins, where (at least if studied from a distance) the image seems to convey more detail about the original image. Here the original values have been binned in such a way that each bin contains approximately the same number of pixels, which means that the entropy of the image is close to the

¹AIBO is a registered trademark of SONY Corporation.

maximum of $\log_2 5 \approx 2.32$. This can also be considered as adding more resolution where most of the data is located.

As discussed in the introduction, it has been found that the contrast encoding in the visual system of the fly is adapted to the specific statistics of its environment [8]. This basically means that, just as in the image of Claude Shannon above, the entropy of the graded response is maximized. More formally, given a sensor X we want to find a partitioning of the data into the N bins of the alphabet $\mathcal{X} = B_1 \cup \dots \cup B_N$ such that each bin B is equally likely. That is,

$$P(X = c \in B_i) \approx \frac{1}{N} \quad (5)$$

which implies that the entropy of which bin data falls into is maximized.

The experiments performed by Brenner et al. [3] also indicate that this kind of entropy maximization constantly is happening in the motion sensitive neurons of the fly. This can be implemented by estimating the probability distribution each time step of the most recent T time steps and changing the transfer function accordingly. In the experiments performed in this paper we have implemented this algorithm using histogram estimators to estimate the probability distributions. In our implementation all instances of the same value are added to the same bin, which explains why the distribution in fig. 1(f) is not completely uniform. The sliding window in our implementation does not use decay, which means that more recent values do not affect the distribution more than older ones within the window.

III. SENSORY RECONSTRUCTION METHOD

In the sensory reconstruction method [15], [13] sensoritopic maps are created that show the informational relationships between sensors, where sensors that are informationally related are close to each other in the maps. The sensoritopic maps might also reflect the real physical relations and positions of sensors. For example, if each pixel of a camera is considered a sensor, it is possible to reconstruct the organization of these sensors even though nothing about their positions is known. It is important to note that using only the sensory reconstruction method, only the positional relations between sensors can be found, and not the real physical orientation of the visual layout. To do this requires higher level feature processing and world knowledge or knowledge about the movement of the agent [13].

To create a sensoritopic map the value for each sensor at each time step is saved, where in this paper each sensor is a specific pixel in an image captured by the robot. A number of frames of sensory data are captured from the robot and each frame is one time step. The first step in the method is to compute the distances between each pair of sensors. This is computed by considering the time series of sensor values from a particular sensor as an information source X . The distance between two sensors X and Y is then computed using the information metric, equation (4).

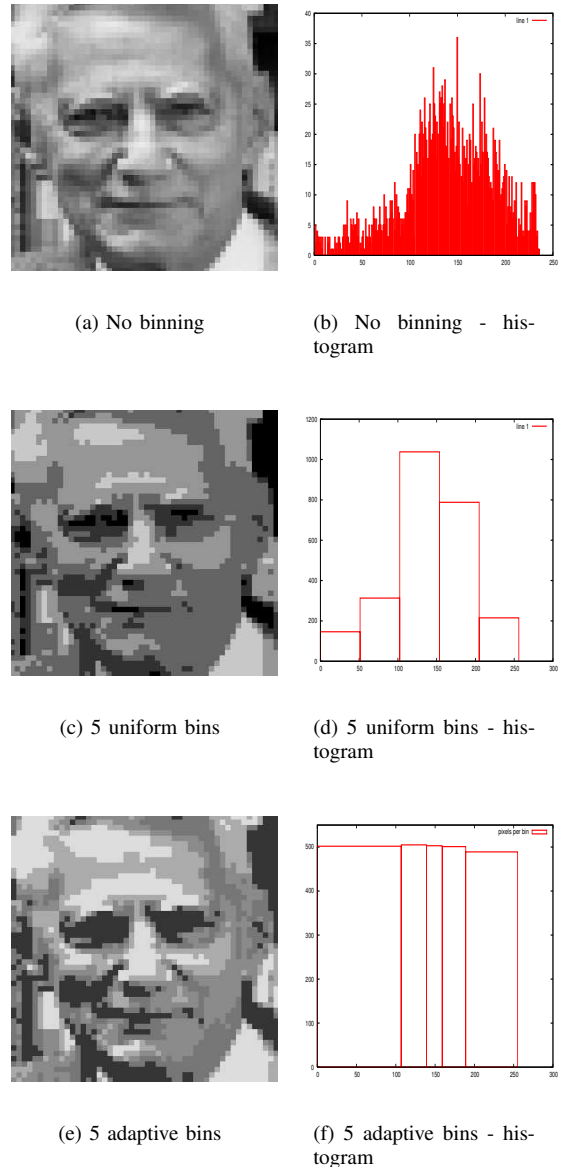


Fig. 1. Example of adaptive binning. fig. 1(a) shows a 50 by 50 pixels grey-scale image of the founder of information theory, Claude Shannon, and fig. 1(b) the corresponding histogram of pixels between 0 and 255. fig. 1(c) shows the same picture where the pixel data (0-255) is binned into only 5 different pixel values using uniform binning and fig. 1(d) the frequency of pixel values. Finally, fig. 1(e) shows the same picture with the pixel data binned into 5 bins using adaptive binning and fig. 1(f) the corresponding histogram. The entropy of the normal binning distribution is ≈ 1.97 while the entropy for the adaptive binning distribution is close to the theoretical maximum of $\log_2 5 \approx 2.32$. The adaptive binning (entropy maximization) increases the resolution where most of the data is located.

From this 2-dimensional distance matrix a 2-dimensional sensoritopic map can be created. In our experiments we have used the relaxation algorithm described in [15].

IV. EXPERIMENTS

A. Simulation

On a 500 by 350 pixel environment (see fig. 2) an 8 by 8 pixel agent represented as a square moves a maximum of one pixel per time step in the x-direction and a maximum

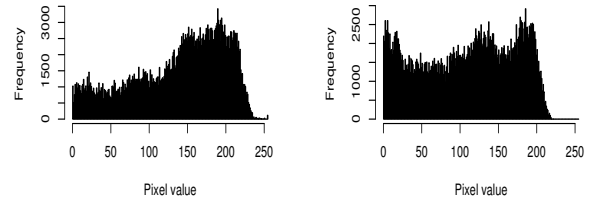
of 1 pixel in the y -direction. Hence dx and $dy \in \{-1, 0, 1\}$, but both cannot be 0 at the same time. Each time step there is a 15% probability that either dx or dy , or both, change value by -1 or 1. Every pixel n ($1 \leq n \leq 64$) of the agent has 4 sensors, one for the red intensity (R_n), one for the green (G_n), one for the blue (B_n), and one for the average intensity of that pixel (I_n). Thus, the agent has a total of 256 sensors. For each time step the values of all the 256 sensors are used as the input to the sensory reconstruction method.



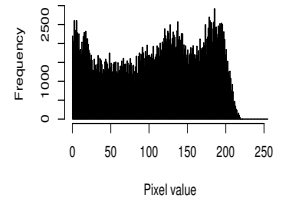
Fig. 2. The environment where the agent is moving. The image depicts autumn leaves and has higher variation in the red and green channels than the blue channel.

Fig. 3 shows the histograms of all sensors of each type accumulated over the whole simulation of 6000 time steps, and also examples of histograms for each sensor type over 10 consecutive time steps. The red and green sensors are quite uniformly distributed over almost the whole range while the blue has a high peak at 0. In fig. 3(d) to 3(f) we can see that the ranges of values in the red and green sensors are greater than in the blue sensors during these 10 time steps, something that was true for most frames.

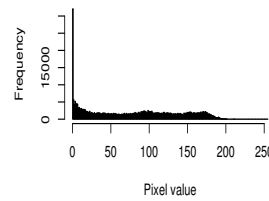
Given the structure of the input data it is expected that adaptive binning with a sliding window would be advantageous for the sensory reconstruction method. In fig. 4(a) the input to the sensory reconstruction method is sensory data from the 256 sensors partitioned into 16 uniform bins (4 bits per sensor). The graph shows that some structure is found and some sensors that are closely positioned in the agent are close in the sensoritopic map. One exception is the blue sensors, B_1 to B_{64} , all located to the left. Clearly, if all the informational structure could be found the map should correspond to the physical order and, for example, R_1 should be close to B_1 . Now consider fig. 4(b). Here the input data was binned into only 4 bins (2 bits per sensor) using entropy maximization with a sliding window of size 100. Here the sensoritopic map clearly shows the informational and physical relationships between the sensors, where sensors that are closely located in the layout of the agent are clustered in the map. This means that the real physical layout of the sensors has been recovered from the raw input data, something that the same method failed to do using uniform binning and double the amount of resolution per sensor.



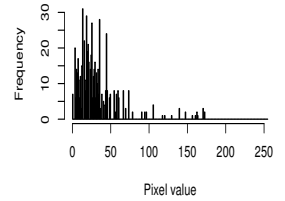
(a) Red histogram 6000 time steps



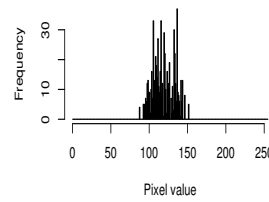
(b) Green histogram 6000 time steps



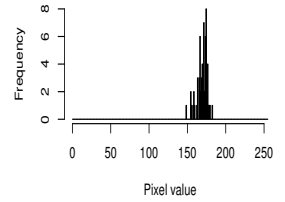
(c) Blue histogram 6000 time steps



(d) Red histogram 10 time steps



(e) Green Histogram 10 time steps



(f) Blue Histogram 10 time steps

Fig. 3. Figures 3(a), 3(b), and 3(c) shows histograms of red, green, and blue sensors from the image in fig. 2 collected from 6000 timesteps of movement from all sensors. Figures 3(d), 3(e), and 3(f) show examples of histograms from 10 consecutive time steps.

B. Robotic Experiments

The robotic experiments were performed with a SONY AIBO robot wandering in an office environment with both artificial lights and two windows. Images of size 88 by 72 pixels from the robot's camera were saved over a dedicated wireless network with an average frame rate of 15 frames per second. The images were transformed to 8 by 8 pixel images by either pixelation with averaging or by using only 8 by 8 pixel from the centre of each image. Either transformation produced similar results in subsequent experiments. Each pixel has three sensors, $R_n, G_n, B_n, 1 \leq n \leq 64$, with R_1 in the upper left corner and R_{64} in lower right corner. The robot performed a simple exploring behaviour

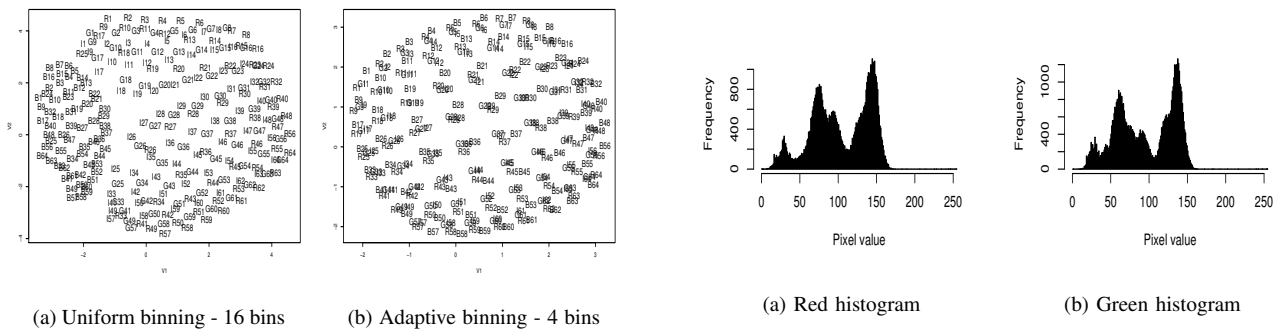


Fig. 4. Sensoritopic maps of the sensory data where R1-64 is the red channel sensors, G1-64 the green, B1-64 the blue, and I1-64 the intensity sensors using uniform and adaptive binning .

with obstacle avoidance walking around in the office.

Fig. 5 shows histograms of all sensors of each type (red, green, blue, intensity) combined over 1000 time steps where each sensor can have a value in the range $\{0, 1, \dots, 255\}$. The red and green sensors have most values between 0 and 170 with two clear peaks at roughly 70 and 150. The blue sensors had a narrower range, with most sensor values between 0 and 80, and two narrow peaks at roughly 25 and 75. The peaks are due to the windows; when walking towards the windows the ambient light is brighter. Similarly to the simulation above, the histograms in any given frame show a narrower range of the data for the blue sensors. Thus, it is expected that the blue sensors are more difficult to reconstruct using uniform binning.

As seen in fig. 6 this is the case. fig. 6(a) shows a sensoritopic map of the blue sensors constructed from 16 uniform bins. Contrast this with fig. 6(b) using only 6 adaptive bins where the organisation of the visual field has been reconstructed. In fig. 7 sensoritopic maps of all the red, green and blue sensors combined are shown. Here we can again see how the adaptive binning enables the sensory reconstruction method to find the positional relations of the sensors of the different types of sensors, where sensor from the same physical position, e.g. R_8 , G_8 , and B_8 , are clustered together. We can also see that the order with R_1 at the opposite corner of R_{64} has been found. In the sensoritopic maps in fig. 7(c) and 7(a) created using 6 and 16 uniform bins we see that the blue sensors are separated from the red and green. The structure of the red and green sensors is also less clear compared to the adaptive binning of fig. 7(b) and 7(d).

V. CONCLUSIONS

This short paper has discussed entropy maximization of sensory data in the fly visual system and how a similar system can be implemented in a robot. The system constantly adapts the input/output mapping of sensory data by estimating the distribution of input data and adapts the output distribution by entropy maximization of the data. This mapping of input/output data compresses the data while maintaining correlations between sensors. Results

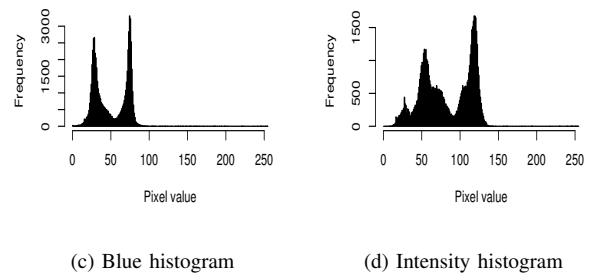


Fig. 5. Histograms of red, green, blue, and intensity sensors.

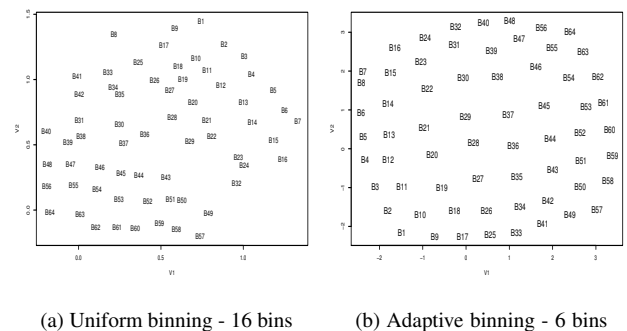


Fig. 6. Sensoritopic maps of blue sensors. fig. 6(a) shows the sensoritopic map created with 16 uniform bins and fig. 6(b) with six adaptive bins using a sliding window.

from simulation show how an agent using this adaptive technique can reconstruct its visual field with a resolution of only two bits per sensor using the sensory reconstruction method. Using four bits per sensor and uniform binning the sensory reconstruction method failed to reconstruct the visual field. This result indicates that adaptive binning is useful for compressing sensory data and to find correlations between sensors. Results from experiments with a SONY AIBO robot show some statistical properties of different indoor environments and how adapting to this structure helps the robot find structure in the sensory data.

Adaptation of sensory systems to the specific environment of a particular species is also studied in the field of sensory ecology [7]. Here many results also seem to

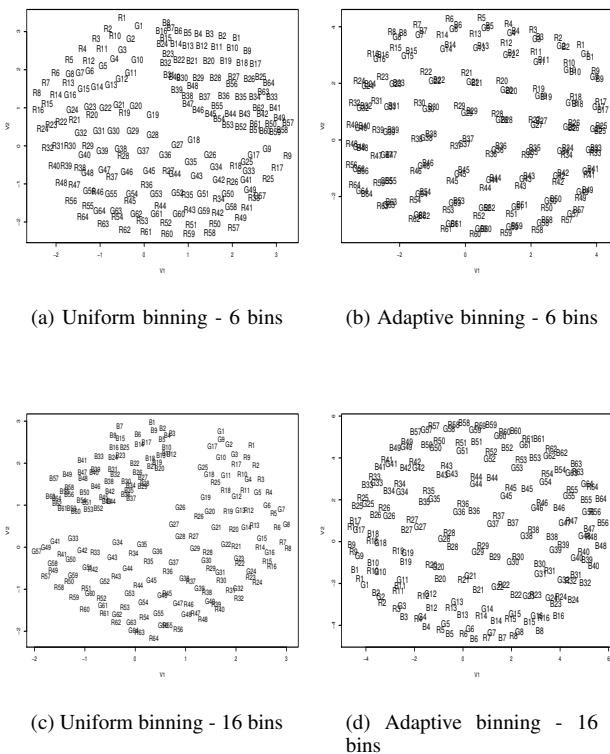


Fig. 7. Combined sensoritopic maps of red (R), green (G), and blue (B) sensors using 6 and 16 bins with uniform and adaptive binning

indicate that sensory systems evolve to be “tuned” to the average statistical structure of the environment. For example, it has been found that the spectral sensitivity of many unrelated fishes has converged to similar patterns depending on the water colour and ambient light.

It is interesting to note how the sensory reconstruction method manages to merge sensors of different types such as red, blue, and green light sensors and reconstruct their visual layout without any knowledge of their physical structure as seen in fig 7(b). This is an example of autonomous sensory fusion. Perhaps the most studied example of this in neuroscience is the optic tectum of the rattlesnake, where nerves from heat-sensitive organs are combined with nerves from the eyes [12]. This kind of multimodal sensor integration is something that will be studied in future work.

This paper has presented initial work and there are many possibilities for developing and extending these ideas. Is entropy maximization more effective for the agent on sets of sensors or a more abstract object level than the single sensor level? In this case, how are these sets or objects selected? There are also many problems associated with entropy estimation. For example, what is the best method to estimate entropy given distributions with empty bins? Here other methods like Miller-Madow bias correction [14] could be investigated. We are also interested in robots that develop and learn over time and adapt to their particular environment and tasks. Results indicate that constraints on

perception during development may improve the perceptual efficiency by reducing the information complexity [19]. In this context adaptive binning can be applied to a developmental system starting with low resolution, where more resolution is added using adaptive binning as the robot develops and learns about its particular environment.

ACKNOWLEDGMENTS

We wish to express our gratitude to the anonymous reviewers for their helpful and insightful comments.

REFERENCES

- [1] H. B. Barlow, “Possible principles underlying the transformation of sensory messages,” *Sensory Communication*, pp. 217–234, 1961.
- [2] W. Bialek, “Physical limits to sensation and perception,” *Ann. Rev. Biophys. Biophys. Chem.*, vol. 16, pp. 455–478, 1987.
- [3] N. Brenner, W. Bialek, and R. de Ruyter van Steveninck, “Adaptive rescaling maximizes information transformation,” *Neuron*, vol. 26, pp. 695–702, 2000.
- [4] E. M. Callaway, “Visual scenes and cortical neurons: What you see is what you get,” *Proceedings of the National Academy of Sciences*, vol. 95, no. 7, pp. 3344–3345, 1998.
- [5] J. P. Crutchfield, “Information and its Metric,” in *Nonlinear Structures in Physical Systems – Pattern Formation, Chaos and Waves*, L. Lam and H. C. Morris, Eds. Springer Verlag, 1990, pp. 119–130. [Online]. Available: citeseer.nj.nec.com/crutchfield90information.html
- [6] R. de Ruyter van Steveninck, J. Anderson, and S. B. Laughlin, “The metabolic cost of neural information,” *Nature Neuroscience*, vol. 1, no. 1, 1996.
- [7] J. A. Endler, “Signals, signal conditions, and the direction of evolution,” *American Naturalist*, vol. 139, pp. 125–153, 1992.
- [8] S. B. Laughlin, “A simple coding procedure enhances a neuron’s information capacity,” *Z. Naturforsch.*, vol. 36c, pp. 910–912, 1981.
- [9] R. Linsker, “Self-organisation in a perceptual network,” *IEEE Computer*, vol. 21, pp. 105–117, 1988.
- [10] J. P. Nadal and N. Parga, “Sensory coding: information maximization and redundancy reduction,” *Neural Information Processing*, vol. 7, pp. 164–171, 1999.
- [11] C. L. Nehaniv, “Meaning for observers and agents,” in *IEEE International Symposium on Intelligent Control / Intelligent Systems and Semiotics, ISIC/ISIS’99*. IEEE Press, 1999, pp. 435–440.
- [12] E. A. Newman and P. H. Hartline, “Integration of visual and infrared information in bimodal neurons of the rattlesnake optic tectum,” *Science*, vol. 213, pp. 789–791, 1981.
- [13] L. Olsson, C. L. Nehaniv, and D. Polani, “Sensory channel grouping and structure from uninterpreted sensor data,” in *2004 NASA/DoD Conference on Evolvable Hardware June 24-26, 2004 Seattle, Washington, USA*. IEEE Computer Society Press, 2004, pp. 153–160.
- [14] L. Paninski, “Estimation of entropy and mutual information,” *Neural Computation*, vol. 15, pp. 1191–1254, 2003.
- [15] D. Pierce and B. Kuipers, “Map learning with uninterpreted sensors and effectors,” *Artificial Intelligence*, vol. 92, pp. 169–229, 1997.
- [16] D. Polani, T. Martinez, and J. Kim, “An information-theoretic approach for the quantification of relevance,” in *Advances in Artificial Life, 6th European Conference, ECAL 2001, Prague, Czech Republic, September 10-14, 2001, Proceedings*, ser. Lecture Notes in Computer Science, J. Kelemen and P. Sosik, Eds. Springer Verlag, 2001, pp. 704–713.
- [17] F. Rieke, D. Warland, R. de Ruyter van Steveninck, and W. Bialek, *Spikes: Exploring the neural code*, 1st ed. MIT Press, 1999.
- [18] C. E. Shannon, “A mathematical theory of communication,” *Bell System Tech. J.*, vol. 27, pp. 379–423, 623–656, 1948.
- [19] A. Slater and S. Johnson, “Visual sensory and perceptual abilities of the newborn: beyond the blooming, buzzing confusion,” in *The development of sensory, motor, and cognitive abilities in early infancy: from sensation to cognition*, F. Simion and G. Butterworth, Eds. Hove: Psychology Press, 1997, pp. 121–141.
- [20] R. Steuer, J. Kurths, C. Daub, W. J., and J. Selbig, “The mutual information: Detecting and evaluating dependencies between variables,” *Bioinformatics*, vol. 18, pp. 231–240, 2002.