

# Evolving Coordinated Behavior by Maximizing Information Structure

Olaf Sporns<sup>1</sup> and Max Lungarella<sup>2</sup>

<sup>1</sup>Department of Psychological and Brain Sciences, Indiana University, Bloomington IN 47405, USA

<sup>2</sup>Department of Mechano-Informatics, University of Tokyo, 113-8656 Tokyo, Japan  
osporns@indiana.edu

## Abstract

Embodied systems actively structure information sampled by their sensors as they engage in sensorimotor interactions with their environment. Can information structure serve as an evolutionary principle that shapes behavior and leads to increased coordination? Here we address this question by attempting to evolve coordinated behavior in a simulated creature subjected to behavioral and information-theoretical cost functions. Our results show that maximizing information structure is highly effective in generating coordinated behavior, providing further support for a potential central role of actively generated information structure in embodied cognition.

## Introduction

Traditional notions of cognition as a computational process that is carried out by disembodied centralized control architectures have been challenged by proponents of embodied cognition (Chiel and Beer, 1997; Clark, 1997; Pfeifer and Scheier, 1999; Sporns, 2002; Pfeifer et al., 2006). Embodied cognition views cognitive processes as arising from a brain embedded in a body with specific perceptual and motor capabilities which interacts with and moves in the real world. From an embodied perspective, coherent, coordinated or intelligent behavior critically depends on the dynamic coupling between nervous system, body, and environment.

An important implication of this view is that embodied systems are not passively exposed to sensory information, but rather actively structure and exploit such information, due to their particular morphology and environmental interactions. It follows that embodiment plays an important and complementary role to neural information processing: the agent's control architecture (e.g. neural system) attends to and processes streams of sensory stimulation, and ultimately generates sequences of motor actions which in turn drive the further selection and production of sensory input. In other words, "information structuring" by sensorimotor activity and "information processing" by the neural system are reciprocally linked.

We have suggested (Lungarella and Pfeifer, 2001; Sporns and Pegors, 2004; Lungarella et al., 2005) that information structure (e.g. statistical regularities, correlations, and redundancies) in the sensory experience of (natural and artificial) embodied systems is – at least in part – the result of self-generated and coordinated motor

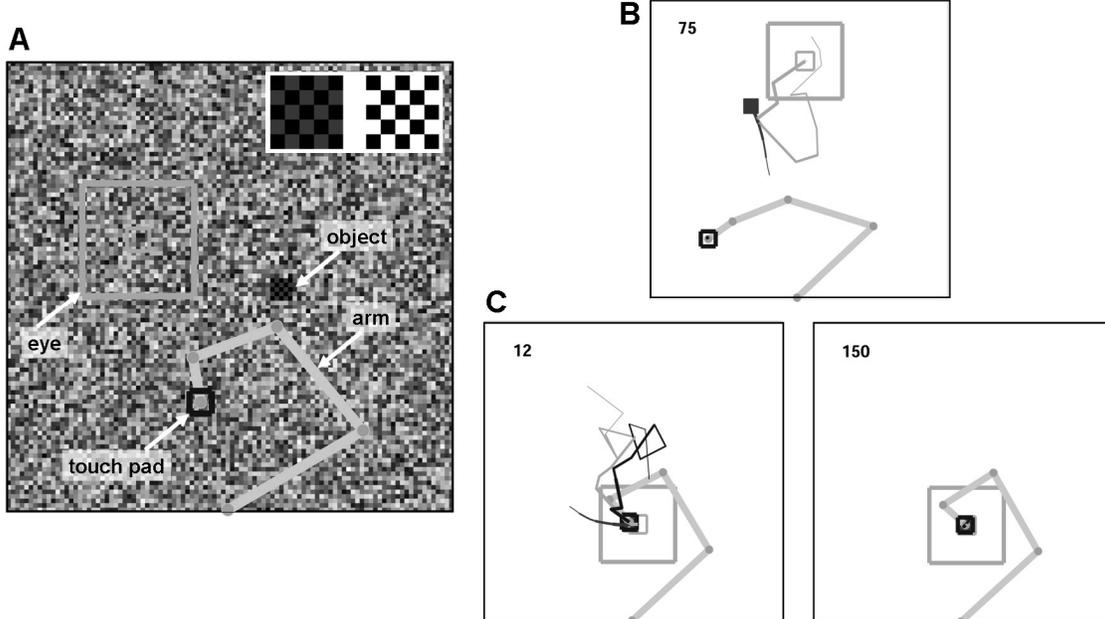
activity. Structured sensory information induced by the interaction with the environment may facilitate neural computation, have a powerful influence on neural learning, and be a central prerequisite for abilities such as multimodal sensory integration, cross-modal learning, and perceptual categorization in the real world.

In previous work (Lungarella et al., 2005), we provided quantitative evidence for the hypothesis that information structure in the sensory experience can be induced by effectively coordinated motor activity. Specifically, we observed a pattern of decreased entropy, and increased mutual information, integration and complexity in sensory data obtained in the course of highly coordinated behavior. Here, we examine the hypothesis that these information-theoretical measures can be used as optimization tools to shape behavior and create or evolve high levels of sensorimotor coordination. More generally, we address the question of whether evolutionary pressure to generate structured sensory information might underlie sensorimotor coordination in behaving organisms.

## Methods

### Sensorimotor Simulation

To investigate the capacity of information-theoretical measures to drive the emergence of coordinated behavior in an optimization, we simulated a simple creature using an evolutionary algorithm implemented in Matlab 7 (Mathworks, Natick, MA). The creature (Fig. 1A), similar in its morphology to DARWIN III (Reeke et al., 1990), performed simple motor acts that resulted in sensory stimulation. Its actions were controlled by a small set of parameters ("genes"), which were varied (mutated and selected) in the course of an optimization process driven by a cost function measuring a specific aspect of performance or "fitness" (see below). The environment was a 2-dimensional bounded area (100×100 pixels), with each pixel at each time step displaying an uncorrelated random color (R, G, and B values chosen with uniform probability from the [0,255] interval). A single object (5×5 pixels) with constant color properties moved in a random path at constant velocity. In addition to color, the object had tactile features (ridges or knobs; Fig. 1A). Tactile features were absent from any other portion of the environment.



**Fig. 1:** Creature and environment. (A) Sample frame showing color distribution in environment, object, eye including fovea, and arm with touch pad. Insert at top right shows visual (left) and tactile features (right) of the object. (B) Sample frame (color environment omitted) showing behavior before evolutionary selection (*random* genomes). (C) Sample frames showing behavior after evolutionary selection for complexity of sensory information (cost function *Cplx*).

The creature possessed two senses, vision and touch. An “eye” (25×25 pixels moving window) was swept across the environment by the action of two (pan/tilt) movement neurons. The center of the eye was designated the fovea – a patch of 5×5 pixels. A touch pad (5×5 tactels) was affixed to the tip of a four-jointed arm (the “hand”) and could be moved across the environment by the action of four joint actuators. If the touch pad made contact with the object, its movement was arrested (the object was “gripped”), while the motion of the touch pad across the surface of the object remained enabled. If touch was lost, the object resumed its random path.

Neural structures of the creature consisted of sensory arrays sampling visual and tactile inputs, a visual attention system, and motor units whose activity controlled eye and arm. The visual attention system allowed the creature to select a salient location within its visual field, on the basis of the activation of a saliency map (Itti et al., 1998; Lungarella et al., 2005). From the raw visual image, we first extracted four opponent color maps to which we applied a winner-take-all mechanism resulting in color feature maps  $I_R$ ,  $I_G$ ,  $I_B$ , and  $I_Y$ . These maps were then linearly combined to form the saliency map:

$$S_{\text{col}} = f(\eta_R I_R + \eta_G I_G + \eta_B I_B + \eta_Y I_Y) \quad [1]$$

with scaling factors  $\eta_R$ ,  $\eta_G$ ,  $\eta_B$  and  $\eta_Y$  quantifying the relative saliency of each of the four color components and  $f$

denoting a sigmoidal activation function. The target eye position  $\theta_E$  coinciding with the instantaneous maximum of the saliency map was used to compute the eye movement

$$M_E = \varphi_E - \theta_E + \rho_E \quad [2]$$

where  $\varphi_E$  denotes the current eye position and  $\rho_E$  is a Gaussian random jitter of the eye.

Prior to all simulation runs, a pre-training procedure was employed to create a mapping between hand positions and joint angle configurations, by randomly actuating the arm across its entire workspace. In the subsequent simulations, target hand locations were then used to retrieve the target angles  $\theta_k$  of the  $k^{\text{th}}$  joint and to calculate the movement command

$$M_k = c_k (\varphi_k - \theta_k) \quad [3]$$

where  $c_k$  is a scaling parameter and  $\varphi_k$  is the current motor command. A Gaussian random jitter  $\rho_4$  was added to the final (fourth) joint angle ( $k=4$ ).

The arm remained in a canonical position, unless the sum of all activation in the foveal region (fov) of the saliency map exceeded a threshold  $\zeta$ :

$$\sum_{\text{fov}} S_{\text{col}} > \zeta. \quad [4]$$

We note that the present paper does not include coupling constants between individual joints and does not address the evolution of joint coordination.

## Evolutionary Algorithm

Eleven parameters (here called “alleles”) collectively controlling the sensorimotor behavior of the creature were subjected to evolutionary optimization (numbers in brackets give initial mean values  $G_i(0)$ , and the initial  $\mu_i(0)$  and final  $\mu_i(G)$  mutation rates (Eq. 5) for the  $i^{\text{th}}$  allele ( $A_i$ ) with  $G$  denoting the total number of generations):

- $A_{1-4}$  (0.50, 0.20, 0.01): These alleles correspond to the four saliency factors,  $\eta_R$ ,  $\eta_G$ ,  $\eta_B$  and  $\eta_Y$  (Eq. 1).
- $A_5$  (0.50, 0.20, 0.01): This allele specifies the amplitude of random jitter movements  $\rho_E$  of the eye (Eq. 2); if set to zero, no jitter occurs.
- $A_{6-9}$  (0.25, 0.20, 0.01): These alleles are movement scale factors  $c_k$  for the four individual joints (Eq. 3), from the “shoulder” at the attachment point (gene 6) to the “fingers” at the tip of the hand (gene 9).
- $A_{10}$  (0.50, 0.20, 0.01): This allele specifies the threshold  $\zeta$  at which arm motion under visual guidance is activated (Eq. 4).
- $A_{11}$  (2.50, 1.00, 0.05): This allele specifies the amplitude of random jitter movements of the terminal joint of the arm  $\rho_4$ , resulting in angular displacement of the attached touch pad (Eq. 3).

At the beginning of each generation  $g$ , a given allele  $A_i$  was mutated as

$$A_i(g) = A_i(g-1) + \gamma(0, \mu_i(g)) \quad [5]$$

the second term denoting a random number drawn from a Gaussian distribution with zero mean and a standard deviation given by the mutation rate  $\mu_i(g)$  for that generation. The mutation rate decreased over time according to

$$\mu_i(g) = \mu_i(G) + (\mu_i(0) - \mu_i(G)) (1 - g/G)^3. \quad [6]$$

Generations consisted of 10 individuals. Each individual was placed in an identical environment and allowed 250 time steps of behavior. At the end of each generation, cost functions were computed for all individuals. The single individual with the maximal value of the cost function was propagated into the next generation, together with 9 mutated offspring. Simulations generally reached stable solutions after  $G=60$  generations.

In all cases, sensory data was collected from the  $5 \times 5$  fovea of the “eye” and from the  $5 \times 5$  touch pad, resulting in sensory time series  $X(t)$  with each variable  $x_i(t)$  corresponding to a single “sensor” (pixel or tactel).

## Cost Functions and Performance Measures

We utilized a variety of cost functions (Table 1), which can be divided into two broad classes: behavioral cost functions and information theoretical cost functions. Behavioral cost functions were based on simple behavioral performance measures such as successful foveation or tactile contact. Information-theoretical cost functions were based on quantifying the statistical structure of sensory inputs sampled by eye and hand. In addition to these cost

functions, we also investigated two “control” cases in which the cost functions were designed to produce “sub-optimal” or “uncoordinated” behavior.

**Table 1:** Behavioral (B), information-theoretical (I) and control (C) cost functions.

Cost Func.		Description
B	<i>foveation</i>	maximizing time for which distance between eye and object is less than 2.5 pixels
	<i>touch</i>	maximizing time for which object is touched
	<i>fovtouch</i>	conjunction of <i>foveation</i> and <i>touch</i>
	<i>maxred</i>	maximizing color red in the fovea (Eq. 7)
I	<i>negH</i>	minimizing entropy
	<i>MI</i>	maximizing mutual information
	<i>Intg</i>	maximizing integration (Eq. 8)
	<i>Cplx</i>	maximizing complexity (Eq. 9)
C	<i>H</i>	maximizing entropy
	<i>negCplx</i>	minimizing complexity (Eq. 9)

Behavioral cost functions were defined as follows. *Foveation* and *touch* maximized, respectively, the ratio between the number of time steps for which the fovea or the touch pad was within 2.5 pixels of the center of the red object ( $\tau_E$ ), and the length of the trial was  $T$  time steps ( $\tau_J$ ). *Fovtouch* was the conjunction of *foveation* and *touch*,  $\tau_{E,J}$ . *Maxred* maximized the amount of the color red, calculated by using the unit activations, integrated over a single trial, within the fovea for the feature maps  $I_R$ ,  $I_G$ ,  $I_B$ , and  $I_Y$ :

$$K_{\text{red}} = I_R / (I_R + I_G + I_B + I_Y). \quad [7]$$

Information-theoretical cost functions and control cost functions evaluated statistical structure of information gathered by the sensors during behavior. These cost functions used measures defined and discussed in detail in a previous publication (Lungarella et al., 2005). We applied four main measures: entropy ( $H$ , *negH*), mutual information ( $MI$ ), integration (*Intg*), and complexity (*Cplx*, *negCplx*). Entropy  $H(X)$  and mutual information  $MI(X)$  were calculated using standard statistical formulations. Mutual information was expressed as the average over all binary relationships between sensors in the sensor data set  $X(t)$ . Integration and complexity (Tononi et al., 1994) capture global aspects of the statistical dependence between elements of a system, as well as the information distribution among such elements. Integration measures the total amount of statistical dependence within a system  $X$  of elements  $x_i$  and can be formulated as:

$$I(X) = \sum_i H(x_i) - H(X). \quad [8]$$

If no statistical relationship between any of the constituent elements  $x_i$  of system  $X$  exists, then  $I(X) = 0$ ; otherwise,  $I(X) > 0$ .

Complexity expresses the degree to which globally integrated and locally segregated information coexists in a data set. It can be expressed as:

$$C(X) = H(X) - \sum_i H(x_i | X - x_i). \quad [9]$$

$C(X)$  is high for (neural) systems that effectively combine functional segregation and integration, for instance, by incorporating specialized elements capable of global (system-wide) interactions. On the other hand,  $C(X)$  is low for random systems, or for systems that are highly uniform, lacking either global integration or local specialization.

Applied to sensory data sets, entropy quantifies the average uncertainty (or self-information) about the state of individual sensors, while mutual information measures the degree to which two sensors are statistically related. Integration serves as the multivariate extension of mutual information, capturing the degree to which all sensors share information. Complexity of sensory data captures the degree to which sensors are specialized (reporting statistically independent events) while at the same time sharing information (as a result of global interdependence).

Quantities used in cost functions ( $\tau_{E,J}$ ,  $K_{red}$ ,  $H(X)$ ,  $MI(X)$ ,  $I(X)$ , and  $C(X)$ ) were also used as performance measures to characterize genomes obtained after optimization (see Table 2).

### Data Analysis

The evolutionary simulations generated sets of genomes (strings of 11 parameters) that gave rise to successful behavior under the cost function used. The performance of the evolved creatures was evaluated by collecting sensorimotor data directly from evolutionary runs or from testing runs performed after evolution was completed. These testing runs consisted of sequences of 100 trials (random initial conditions,  $T=250$ ) during which evolution was “frozen.” The results from the testing runs were used to rank genomes in terms of their performance. The ten best genomes for each cost function were then further analyzed. Standard t-tests were applied to evaluate significant differences ( $p < 0.01$  for all comparisons labeled “significant” in this paper).

In addition to genomes optimized under a cost function we also evaluated the performance of several “constructs” (designed or manipulated genomes). Constructs included a sample of 40 random genomes collected before evolution (*random*), a set of 40 genomes obtained after optimizing for  $Cplx$ , but with  $A_{11}$  set to zero (*nojitter*) and a genome with allele values set to  $[1\ 0\ 0\ 0\ 0\ 1\ 1\ 1\ 1\ 0\ 0]$  (*ideal*).

## Results

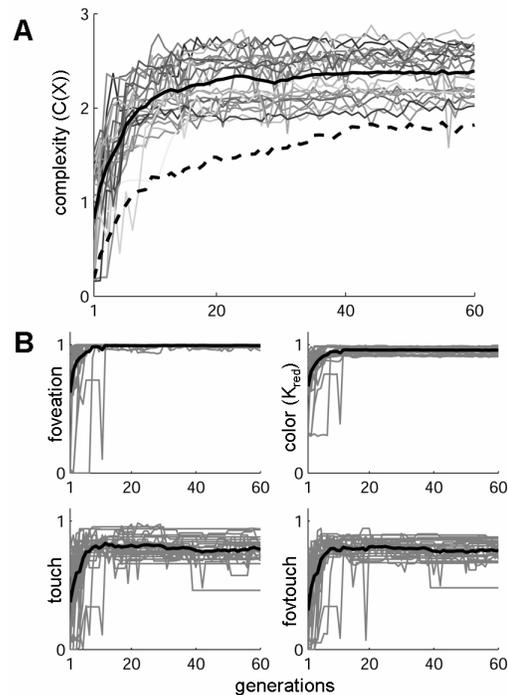
### Evolving Visuomotor Coordination

Figure 1 shows frames from a single behavioral trial, obtained before any evolution has taken place (Fig. 1B), as well as after evolution (Fig. 1C). Before evolution, the arm was stored in a canonical “resting” position, and it activated only if visual activity in the center of the saliency map exceeded the threshold set by  $A_{10}$ . Eye movements

were driven by colors in the environment, with all colors contributing to the generation of the motor command displacing the window ( $A_{1-4}$ ). Occasional contact with the randomly moving object did not result in persistent foveation and tracking movements.

After evolution (cost function  $Cplx$ ; Fig. 1C), the arm tended always to be activated and tracked movements of the eye as it scanned the environment. Once a red object was acquired, eye movements controlled by the activity of the saliency map lead to foveation and tracking. The proximity of the tip of the arm to the visual field facilitated rapid reaching movements establishing contact between the touch pad and the surface of the object. Visual and tactile exploration of the object continued for the remainder of the trial, involving small amplitude visual and tactile scanning movements.

Figure 2 shows the evolutionary time course of the cost function  $Cplx$  for a set of evolutionary optimizations, as well as several simultaneously recorded behavioral indices. The complexity of sensory information increased in accordance with the cost function, as did several behavioral measures (e.g. foveation, touch, and amount of color red). The example shows that maximizing information structure leads to the emergence of visual foveation, tracking, reaching and tactile exploration, yielding a coordinated behavioral pattern.



**Fig. 2:** (A) Time course of cost function  $Cplx$  for most fit individual in 40 optimization runs (grey lines), average of most fit (black line), and average over entire generation (black hatched line). (B) Corresponding behavioral performance measures.

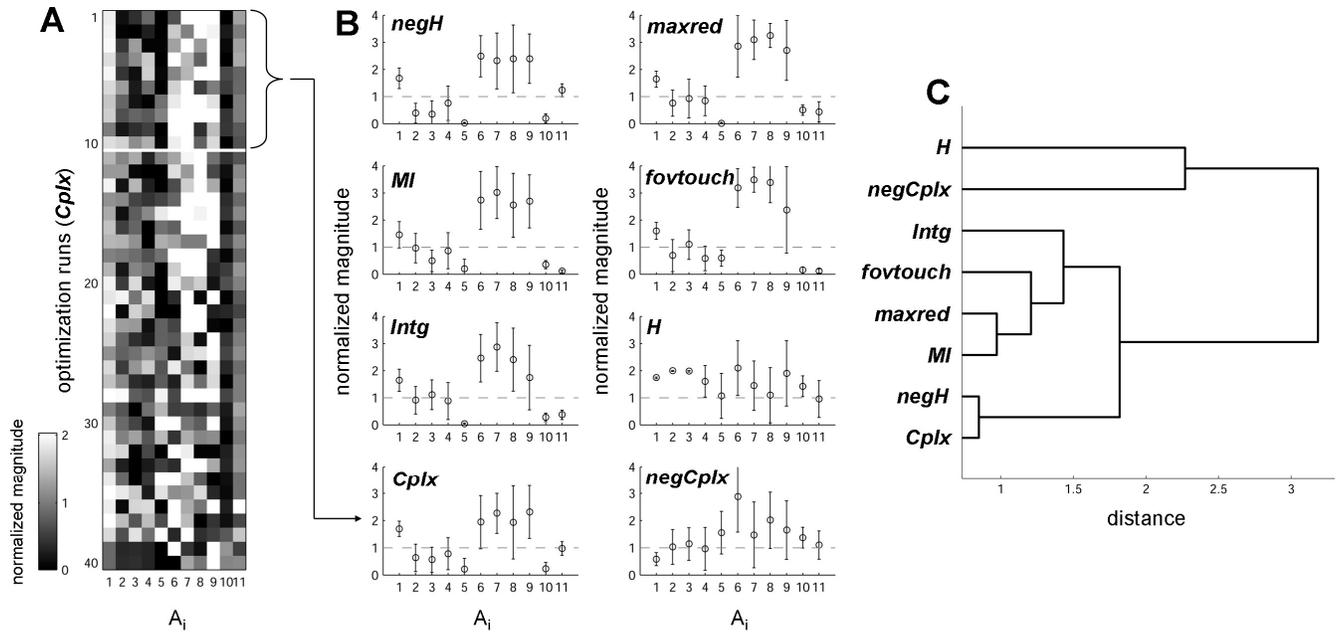
**Table 2:** Performance of the evolved creatures (n=10) for selected cost functions (numbers are means and standard deviations).

Construct	Cost Function	Performance Measure					
		H(X)	MI(X)	I(X)	C(X)	K <sub>red</sub>	τ <sub>E,J</sub>
<i>random</i>		3.52 (0.63)	0.04 (0.04)	23.77 (18.17)	0.30 (0.29)	0.20 (0.28)	0.05 (0.14)
	<i>negH</i>	1.02 (0.07)	0.17 (0.01)	96.16 (11.40)	1.49 (0.07)	0.80 (0.03)	0.68 (0.03)
	<i>MI</i>	1.77 (0.19)	0.34 (0.04)	111.06 (9.63)	0.69 (0.08)	0.77 (0.02)	0.72 (0.08)
	<i>Intg</i>	1.40 (0.12)	0.28 (0.04)	120.56 (4.32)	1.10 (0.14)	0.79 (0.02)	0.71 (0.04)
	<i>Cplx</i>	1.25 (0.23)	0.17 (0.02)	92.24 (7.77)	1.39 (0.04)	0.78 (0.02)	0.64 (0.03)
	<i>maxred</i>	1.38 (0.23)	0.30 (0.08)	122.92 (5.37)	0.99 (0.33)	0.81 (0.02)	0.77 (0.07)
	<i>fovtouch</i>	2.02 (0.30)	0.33 (0.04)	106.88 (15.13)	0.67 (0.17)	0.77 (0.02)	0.79 (0.06)
	<i>H</i>	3.65 (0.19)	0.02 (0.01)	13.40 (1.19)	0.16 (0.01)	0.00 (0.00)	0.00 (0.00)
<i>nojitter</i>		1.94 (0.21)	0.24 (0.06)	85.68 (10.74)	0.70 (0.03)	0.74 (0.12)	0.56 (0.17)
<i>ideal</i>		2.01 (0.05)	0.31 (0.01)	79.46 (2.64)	0.47 (0.02)	0.75 (0.03)	0.96 (0.01)

### Analysis of Evolved Genomes

Figure 3A illustrates the distribution of genomes obtained from 40 evolutionary simulations using the *Cplx* cost function. We found a significant degree of variability across multiple solutions, indicating that the specific values of some genes may vary considerably and yet produce successful behavior. Genomes (ten best from each condition) obtained using several information-theoretical, behavioral and control cost functions are summarized in Fig. 3B. Genomes evolved under information-theoretical and behavioral cost functions are characterized by some overall similarities. For example, all show increased values

for  $A_1$  but decreased values for  $A_{2-4}$ , indicating a preference for the color red during visual behavior. All expressed increases in  $A_{6-9}$ , controlling arm movements, consistent with progressively increased movement amplitudes after evolution. In most cases, jitter of the eye was largely eliminated.  $A_{10}$  was decreased in all cases, indicating that the arm remained activated throughout all behavioral trials. This overall pattern was consistent with all cost functions selecting for visual foveation and tracking of red objects, as well as reaching movements towards the object.



**Fig. 3:** (A) Genomic profiles (normalized to initial mean values) for 40 evolved creatures (*Cplx*), ranked top to bottom after 100 trial testing runs. The ten best genomes were used for further analysis. (B) Mean and standard deviation for ten best genomes (11 alleles) obtained using cost functions *negH*, *MI*, *Intg*, *Cplx*, *maxred*, *fovtouch*, *H* and *negCplx*. (C) Hierarchical cluster analysis and resulting dendrogram, based on average group distances of the 8 mean genomes obtained from data shown in (B).

$A_{11}$  was decreased for cost functions *MI*, *Intg*, *maxred* and *fovtouch*, but remained elevated for *negH* and *Cplx*. A decrease signals elimination of tactile jitter, while an increase indicates persistent exploratory motion by the touch pad on the surface of the object. These exploratory movements had an important role to play in generating additional structure in tactile sensory data.

Control cost functions *H* and *negCplx* showed little systematic difference to the initial random state. Most significant were the increased values for all color-related genes in *H*, as well as for jitter-related genes in both *H* and *negCplx*, consistent with behavior that minimized the amount and structure of sensory information.

To summarize patterns of similarity among the average genomic profiles obtained by using the eight cost functions *negH*, *MI*, *Intg*, *Cplx*, *maxred*, *fovtouch*, *H*, and *negCplx*, we performed a hierarchical cluster analysis using a linkage algorithm based on average unweighted group distances (Fig. 3C). We observe that the control cost functions are well segregated, and that *H* and *Cplx* form a subcluster within the remaining group. The behavioral cost function *maxred* and the information-theoretical cost function *MI* appear highly similar.

We investigated the impact of eliminating or hand-specifying individual genes on behavioral performance and information structure. This strategy can reveal the individual contribution (or lack thereof) of an individual gene on overall behavior. Table 2 shows two such cases, called *nojitter* and *ideal*. Case *nojitter* was obtained by modifying *Cplx* genomes, specifically by setting  $A_{11}$  to zero, thus eliminating touch pad jitter. This resulted in a small and non-significant decrement in behavioral performance (compared to *Cplx*) while it produced a significant decrease in complexity and an increase in entropy. Case *ideal* consisted of a genome constructed on the basis of the expected roles of the individual genes in behavior. The genes were set to the string [1 0 0 0 0 1 1 1 1 0 0], resulting in near perfect foveation and reaching behavior, even exceeding the performance of *maxred* or *fovtouch*. However, a significant amount of information structure was lost, especially as measured by entropy, complexity, and integration.

## Discussion

In this paper, we addressed the question of whether information theoretic cost functions can be used to evolve coordinated sensorimotor behavior. We showed that maximizing information structure is highly effective in producing coordinated behavior in a simple sensorimotor creature. The resulting behavior closely resembled that obtained when using behavioral cost functions that directly evaluate behavioral success or error. Cost functions previously found to be associated with increased coordination (*negH*, *MI*, *Intg*, and *Cplx*; Lungarella et al.,

2005) lead to similar coordinated behavioral solutions, while *H* or *negCplx* produced uncoordinated behavior. There were subtle but significant differences in the behavioral strategies that emerged, for example when comparing *MI* to *negH* or *Cplx*. Maximization of mutual information in sensory inputs tended to suppress active exploration as it favored persistent and uniformly high correlations between sensors, while continued scanning of visual and tactile surfaces generated additional information structure, as captured by increases in complexity.

The focus in this paper is on the structure of information as sampled by the sensory system, but can be readily extended to include informational measures obtained from neural representations. Previous work demonstrated that information structure can drive the emergence of specific neural connectivity patterns and promote matching between sensory inputs and neural systems (Sporns et al., 2000). A related approach involves optimizing behavioral cost functions, and then examining informational measures such as complexity in the control (e.g. neural) structure. Using this approach, Seth and Edelman (2004) found that complex neural dynamics may facilitate adaptation to rich sensory environments and motor demands. Yeager and Sporns (2006) found that the neural networks of organisms evolving in a computational ecology exhibit significant structural elaboration and increases in neural complexity as evolution progresses. Viewed in the context of these earlier studies, our present results suggest that the functional potential of neural architectures may only become fully realized if their neural complexity is matched by rich and highly structured sensory data.

Our creature's vision system and its arm do not realistically model an animal or a robot – does this undermine our general conclusions about embodied systems? Following an argument suggested by Beer (2003) in his study on evolving categorical perception in a simple visually-guided agent, if our goal is to understand how information structure can be used as a means to evolve highly coordinated behavior, physical realism is not required. In a sense, our creature is engaged in a minimally cognitive behavior, that is, behavior which raises cognitively relevant issues (Beer, 2003).

In future work, we will subject additional parameters to evolutionary pressure. Given a particular task environment, various combinations of control system and morphology could be evaluated in terms of their ability to generate informational structure in the sensory channels. It is also possible to envision a scenario in which – based on how efficiently the task is performed – the methods themselves are subjected to co-evolutionary optimization.

## Acknowledgements

M.L. was supported by research grant no. 04413 of the Japanese Society for the Promotion of Science.

## References

- Beer, R.D. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11(4):209-243.
- Clark, A. (1997). *Being There: Putting Brain, Body and World Together Again*. Cambridge, MA: MIT Press.
- Chiel, H.J. and Beer, R.D. (1997). The brain has a body: adaptive behavior emerges from interactions of nervous system, body, and environment. *Trends in Neurosciences*, 20:553-557.
- Itti, L., Koch, C. and Niebur, E. (1998). A model of saliency-based visual attention for rapid scene analysis. *Trans. Pattern Analysis and Machine Intelligence*, 20(11):1254-1259
- Lungarella, M. and Pfeifer, R. (2001). Robots as cognitive tools: Information-theoretic analysis of sensory-motor data. *Proc. of 1<sup>st</sup> Int. Conf. on Humanoid Robots*, pp.245-252.
- Lungarella, M., Pegors, T., Bulwinkle, D. and Sporns, O. (2005). Methods for quantifying the information structure of sensory and motor data. *Neuroinformatics*, 3(3):243-262.
- Pfeifer, R. and Scheier, C. (1999). *Understanding Intelligence*. Cambridge, MA: MIT Press.
- Pfeifer, R., Lungarella, M., Sporns, O. and Kuniyoshi, Y. (2006). On the information-theoretic implications of embodiment – principles and methods. *Trends in Cognitive Sciences* (in press).
- Reeke, G.N., Jr., Sporns, O. and Edelman, G.M. (1990). Synthetic neural modeling: The "Darwin" series of recognition automata. *Proc. IEEE*, 78:1498-1530.
- Seth, A. and Edelman, G.M. (2004). Environment and behavior influence the complexity of evolved neural networks. *Adaptive Behavior*, 12(1):5-20.
- Sporns, O. (2002). Embodied Cognition. In M. Arbib (ed.) *MIT Handbook of Brain Theory and Neural Networks*, Cambridge, MA: MIT Press, pp. 395-398.
- Sporns, O., Tononi, G. and Edelman, G.M. (2000). Theoretical neuroanatomy: Relating anatomical and functional connectivity in graphs and cortical connection matrices. *Cerebral Cortex*, 10:127-141.
- Sporns, O. and Pegors, T.K. (2004). Information-theoretic aspects of embodied artificial intelligence. In F. Iida, R. Pfeifer, L. Steels and Y. Kuniyoshi (eds.) *Embodied Artificial Intelligence*, LNCS 3139. Heidelberg: Springer-Verlag, pp.74-85.
- Tononi, G., Sporns, O. and Edelman, G.M. (1994). A measure for brain complexity: Relating functional segregation and integration in the nervous system. *Proc. Nat. Acad. Sci. USA*, 91:5033-5037.
- Yaeger, L., and Sporns, O. (2006) Evolution of neural structure and complexity in a computational ecology. In Rocha, L. et al. (eds.) *Artificial Life X*. Cambridge, MA: MIT Press (this volume).