Sensorimotor Information flow in Genetic Regulatory Network driven control systems

Tom Quick[1,2], Chrystopher Nehaniv[2], Kerstin Dautenhahn[2] and Graham Roberts[1]
[1]Adaptive Systems Research Group,
University of Hertfordshire, Hatfield, Herts AL10 9AB, UK
[2]Department of Computer Science,
University College London, Gower St., London WC1E 6BT, UK
{t.quick,g.roberts}@cs.ucl.ac.uk
{k.dautenhahn,c.nehaniv}@herts.ac.uk

Abstract

We present results from applying information theory based measures to simulated robots controlled by evolved Genetic Regulatory Networks. Measuring the information flow across sensory and effector surfaces, we create an information profile illustrated using area-proportional Venn diagrams. We examine the relationship between this information profile and various elements of the overall system including the GRN controlling the robot, the environment, the nature of the task for which the GRN controller has been evolved, and evolutionary fitness.

Introduction

In the natural world, Genetic Regulatory Networks (GRNs) exist as part of a system of interacting elements that shapes how a cell engages with its external environment—typically other cells in the case of eukaryotes or an environment external to the organism in the case of prokaryotes (West-Eberhard, 2003; Alberts et al., 2002). This is a process that continues as an essential part of an organism's development throughout its lifetime (Davidson, 2001; Arthur, 2000).

In Artificial Life research, GRNs have traditionally been studied as stand-alone dynamical systems (Kauffman, 1969; Reil, 1999; de Jong et al., 2003) or as models of developmental processes as discussed in (Dellaert and Beer, 1996). This latter approach has been used to grow neural networks (Astor and Adami, 2000), body plans and nervous systems for artificial agents (Bongard, 2002; Hornby and Pollack, 2002) and to generate other morphologies and pattern formation phenomena (Eggenberger, 1997; Kumar and Bentley, 2003), with varying degrees of biological realism. A common feature of such research is the use of GRNs as a tool for implementing a biologically inspired non-linear mapping from genotype to phenotype. The phenotype is typically matured in isolation from its environment, at which point the GRN is discarded. In contrast, building on previous work (Quick et al., 2003), we use our GRNs as the primary driver of behaviour throughout each simulated agent's lifetime, during which it is constantly engaged with an environment.

In this paper we show and evaluate results using measures from information theory (Shannon, 1948; Cover and Thomas, 2001) represented using area-proportional Venn diagrams (Chow and Ruskey, 2004) as an analytic tool, applied to this relationship.

There is excellent research available on the use of information theory in robotics. Lungarella has produced extensive studies on various measures applied to sensorimotor coordination during foveation tasks (Lungarella et al., 2005). Olsson provides elegant illustrations of information-based measures and metrics applied to robot sensory data sets (Olsson et al., 2004). However, the former focuses on a very specific task were we address a more general and basic case of robot–environment interaction during an evolutionary process, whilst the latter is not concerned with task performance.

In the following section we present our GRN-driven model. We then describe the task for which we evolve solutions: controlling a simple simulated robot to seek light and avoid obstacles. We present our methods of analysis, based on information theory and area-proportional Venn diagrams, and then show our experimental results and apply these methods to them. Finally, we discuss some of the implications from the results and analysis and draw some conclusions.

Biosys: An evolvable GRN-driven control system

Biosys is a biologically derived system that models simple gene-protein-environment interaction. For the experiments described here, the system was set-up as follows.

Each gene has two regulatory sites (*cis*), each represented by an integer in the range 0-7 and a bit indicating whether protein binding at that site has an activatory or inhibitory effect on the gene's expression.

Proteins, also represented by integers in the same range, acting as transcription factors (*trans*) bind to *cis* sites that have identical integer values to themselves. Each gene outputs instances of a single protein, again in the range 0-7, henceforth denoted p*n*. The number of instances output is

determined by summing the number of proteins bound to each regulatory site. An overall activation level for each gene is derived by adding the figures for activatory *cis* sites and subtracting for inhibitory *cis* sites.

Each gene's activation level is then input into a simple *tanh*-derived output function to determine the number of protein instances produced.

Each gene has a bit determining if it is on or off by default. Genes that are on by default, or 'default active', have a small integer added to their output such that it is > 0 even when their activation level is 0. Negative activation levels make it possible for a gene to be completely suppressed even though it may be on by default.

Proteins that are not bound to *cis* sites can engage with the cell's environment. The system maintains a fixed one-to-one mapping between sensory events and specific protein types, and between specific protein types and actuator events. The challenge for the evolutionary algorithm, described below, is then to evolve GRNs that engage with these and the other non- environmentally mapped proteins over time in such a way that a 'fit' behaviour is produced.

In the case of the Khepera robot used for this paper (simulated using the *YAKS* software package[1]), light from the left side of the robot's body was summed and input as instances of p0, and light from the right as p1. Summed proximity data from the left Infrared sensors was input as p3, and from the right as p4. The light and protein sums were each multiplied by a gain value to determine the number of protein instances to input. A gain of 0.3 was applied to summed light data, and 0.08 to summed proximity data.

The robot's left wheel speed was determined by the level of p6, and the right by p7. These were mapped by calculating each protein type's level as a proportion of a protein saturation value (itself a genome-wide evolvable parameter), and setting the relevant wheel's speed to be the equivalent proportion of its maximum possible value (a feature of the robot or simulation implementation).

This protein-based environmental mapping allows a round-trip from environmental events, to cellular protein levels, to gene activation levels, in turn impacting regulatory dynamics, cellular protein levels, and activity within the environment in an ongoing dynamical process that shapes the system's responses to its environment over time.

Each cell's lifetime is implemented as a series of simulated time-steps, each broken down into sequential bind-transcribe-decay operations. During binding, proteins not already bound to *cis* sites bind to those that they match to, based on simple identity. Where possible, multiple *cis* sites with the same value share an equal number of binding *trans* proteins. The number of proteins that bind as a proportion of the total available is determined by a global genetic parameter that sets the proportion of available matching proteins that bind to *cis* sites.

Next, transcription is performed. Each gene in turn expresses a number of proteins based on its activation level, as described above.

Next, proteins are decayed. This applies whether or not they are bound to *cis* sites. Each protein type has its own decay rate. This specifies the proportion of proteins of that type that decay at each time step. Each decay rate is a global evolvable parameter.

Finally, environmental interaction occurs as described above. Proteins that are involved in environmental output (i.e. proteins 6 and 7, which determine the robot's wheel speed) are automatically decayed, or 'burnt'.

Evolving light-seeking with obstacle avoidance

Individuals, each with 5 genes with 2 *cis* sites on each gene were evolved to maximise their exposure to light during their lifetime in a simple bounded square environment with a light source in the centre.

We tried at first to evolve from a random population in an environment containing an intermittently broken wall (figure 1*(c)*), thus creating both initial exposure to light and the need for obstacle avoidance to achieve greater light exposure, each individual starting their 1000 time-step life positioned behind the wall and oriented at 90 degrees to the light source. Note that we were more interested in investigating the system's ability to evolve a GRN integrating both types of sensor data than potential robustness to environmental variation, so did not vary starting position and orientation.

We used a basic fixed-length haploid Genetic Algorithm (GA) (Goldberg, 1989), with single point cross-over, tournament selection and weak elitism (a single fittest individual being retained from each population into the next). Cross-over occurred with a probability of 0.9, and mutation with a probability of 0.01 per bit. Each population comprised 150 individuals and was evolved over 100 generations for each evolutionary run.

However, fit individuals repeatedly failed to evolve. We believe that this is because the fitness landscape was littered with local minima caused by spaces between wall sections allowing an increase in light reaching the robot's light sensors. Using the proximity sensors and moving away from the wall involved a corresponding drop in exposure to light, and hence fitness.

We therefore tried incrementally evolving in three phases, as illustrated in figure 1. The first phase (a) involved an environment with small solid walls behind which the robot was positioned. The final generation from the first phase was then used as the seed population for the second phase (b), in which the length of the walls was increased. The process was repeated for the third phase (c), in which 'holes' were made in the walls. This incremental regime succeeded in producing individuals able to use their proximity sensors to

[1] http://r2d2.ida.his.se/

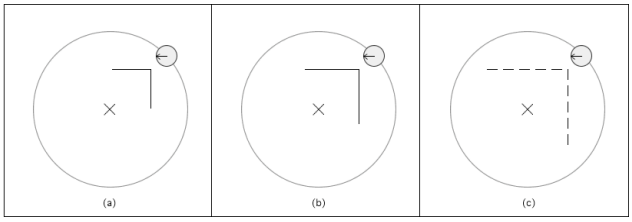

Figure 1: Environments in which light-maximising individuals were evolved incrementally, first for *(a)*, then *(b)* and finally *(c)*. *X* indicates a central light-source, the large circle its diffusion boundary. The small circle denotes each individual's starting position and orientation.

navigate past the walls, entering the region of greater light intensity nearer the centre of their environment.

Methods

This section describes the primary information-theoretic tools and representations used to analyse the results data generated over the three incremental evolutionary phases described above.

Measuring information

Let $\mathcal{X}$ be the alphabet of values of a discrete random variable (information source) X, where X in the present article will represent a sensor or actuator. The probability that the random variable X assumes a value $x \in \mathcal{X}$ is denoted $p(x)$. Then the *entropy*, or uncertainty associated with X is

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

which specifies in bits the information necessary to specify the value assumed by X. The *joint entropy* of two variables X and Y is calculated taking into account the joint probability of every unique combination of the two:

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

The *conditional entropy*,

$$H(Y|X) := \sum_{x\in\mathcal{X}} p(x)H(Y|X=x) \tag{3}$$

$$:= -\sum_{x\in\mathcal{X}} p(x)\sum_{y\in\mathcal{Y}} p(y|x)\log_2 p(y|x) \tag{4}$$

specifies the uncertainty associated with the discrete random variable Y if the value of X is known. This describes how much more information is required to fully determine Y once X is known, or equally, how much information about Y is independent of X. The *mutual information* is the information shared between the two random variables X and Y and is defined as

$$I(X;Y) = H(X) - H(X|Y) = H(Y) - H(Y|X). \tag{5}$$

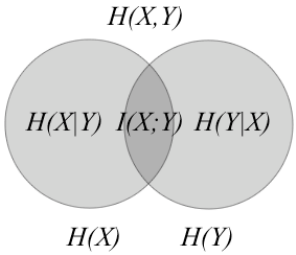

Figure 2: Venn diagram representation of the relations between entropy $H(X)$, conditional information $H(X|Y)$, mutual information $I(X;Y)$ and joint entropy $H(X,Y)$

Area-proportional Venn diagrams

A Venn diagram can be used to illustrate the informational relationship between between mutual information, $I(X;Y)$, entropy, $H(X)$, and conditional entropy, $H(X|Y)$, as portions of the joint entropy $H(X,Y)$ as illustrated in figure 2.

Furthermore, it is possible to generate *area-proportional* Venn diagrams (Chow and Ruskey, 2004). Applying this technique to the information relations as described above, we can create a visual representation of the balance between information flowing into and out of a system via its sensorimotor surfaces.

A single Venn diagram is used to represent the information profile or balance of an individual over the duration of its lifetime (1000 time-steps), such that variation in this balance between different fit individuals from different populations can be readily compared, and evolutionary shifts in GRN dynamics producing different behavioural strategies can be identified in terms of the consequences for the flow of information into and out of the system.

Here we use the region $H(X)$ to represent the joint entropy of the four sets of sensor data that are passed into the cell model as protein instances: light and proximity from the left and right hand side of the robot. $H(Y)$ represents the joint entropy of the two sets of wheel data, left and right.

In calculating $H(X)$, each unique combination of values occurring in the four sensor data sets is treated as a member of the alphabet $\mathcal{X}$, and similarly for calculating $H(Y)$ from the two wheel data sets. $H(X,Y)$ therefore shows the total sensorimotor information of an individual over its lifetime.

The *relative* areas of $H(X)$, sensory information, and $H(Y)$, motor information, illustrate the proportion of sensor to motor information. The relative area of the region $H(X|Y)$ shows the proportion of sensor information that is disregarded insofar as it is not reflected in actuator information, $H(Y)$, and vice-versa for $H(Y|X)$, showing the proportion of actuator information that is independent of sensory information. The relative area of $I(X;Y)$ shows the proportion of information common to both sensors and motors.

In determining the relative areas of $H(X)$, $H(X|Y)$ and $I(X;Y)$ for the area-proportional diagrams, we divide each

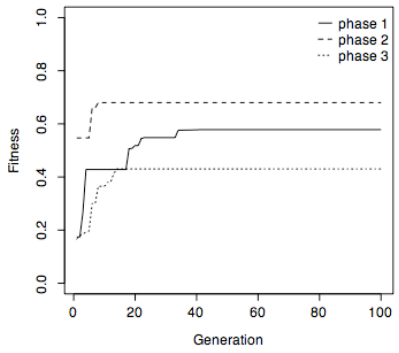

Figure 3: Fitness over 3 phases of incremental evolution. Phases 1–3 correspond to the environments illustrated in figure 1(a)–(c)

of these values by $H(X,Y)$.

Results and Analysis

Fitness over the 3 phases of incremental evolution described above is shown in figure 3. Incremental evolution was successful in so far as individuals were evolved in phase 3 that successfully navigated past the broken wall into the central and well-lit region of the environment. Evolution in phase 2 successfully continued increasing fitness from phase 1 with no significant initial drop in fitness.

Figure 4 shows for phase 1, at points of significant shifts in fitness, an area-proportional information relationship Venn diagram and robot trajectory for the individual lifetime that produced the corresponding fitness measure. Note that the light grey box inside each Venn diagram shows how big the diagram would be if it were scaled to match the actual joint entropy of the data used as a proportion of the maximum possible joint entropy.

Figure 4(a) shows an individual from generation 2. It moves very slowly a short distance to the wall, its left wheel activated weakly by part of the output from a default active gene, and its right by a gene responding to light from the robot's left hand side[2]. This lack of activity is reflected in a low overall value for $H(X,Y)$, hence the grey scaling box is very small. There is very little uncertainty regarding motor output given sensor input: $H(Y|X)$ proportional to $H(X,Y)$ is very small. As with all lifetimes in all phases, the sensor data contains more information than the motor data, $H(X) > H(Y)$, a observation which is discussed below.

In generation 3, the relationship with the environment changes dramatically due to a small genetic network wiring shift and a small global parameter change that enables the system to engage with its environment at a basic level. Left hand side light activates the right wheel while the left turns at a constant rate such that the robot is able to move parallel

[2]The YAKS simulator models low height walls, allowing some light to reach the robot when it is behind a wall

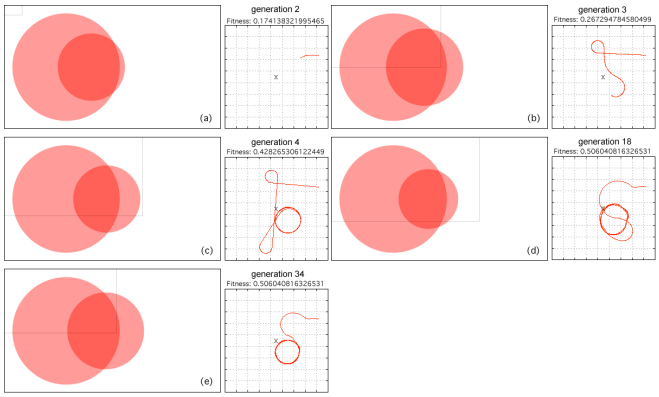

Figure 4: Area-proportional information relationship Venn diagrams and robot trajectories from individuals at key fitness points in evolutionary phase 1, in environment shown in figure 1(a). The light-grey box inside each Venn shows how big the diagram would be if it were scaled to match the joint entropy present as a proportion of the maximum possible joint entropy.

to the wall. When it exits the lit area the right wheel slows, causing the robot to loop to the right, until it re-enters the light and resumes a trajectory towards the light source.

Generation 4 (figure 4(c)) sees a small alteration to maximum protein levels resulting in an almost identical trajectory but at a faster speed, producing a major increase in fitness. This increase in activity produces a greater amount of information, $H(X,Y)$, and a slight increase in sensory information relative to actuator information.

By generation 18, GRN re-wiring has produced an individual sensitive to left hand side proximity data, enabling it to move much more directly toward the light source by 'knowingly' dodging the wall, although the rest of the behavioural strategy is similar to generation 4. $H(X,Y)$ is further increased, whilst the proportions of sensory information increases relative to motor information.

Generation 34 sees another increase in GRN complexity in terms of the degree of interconnection between genes, with both left and right proximity data being used, and multiple cyclic connections in inter-gene wiring. Despite this, $H(X,Y)$ decreases, perhaps due to the more regular trajectory resulting in more repetitive and hence probable sensor and motor data . $H(Y|X)$ reaches its probable peak for this evolutionary phase, possibly reflecting the information generated by the GRN's dynamics independently of sensory input. $H(Y)$ has also increased relative to $H(X)$, perhaps also due the regularity of the trajectory and hence similar information levels in sensor and motor data given this fairly featureless environment.

The major evolutionary features in phase 2 are shown in figure 5. Individuals from the final generation of phase 1 were already adapted to this new environment with longer

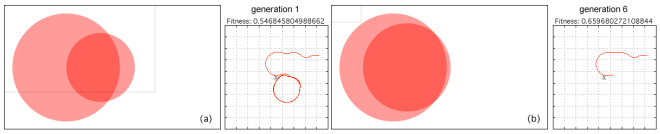

Figure 5: Area-proportional information relationship Venn diagrams and robot trajectories from individuals at key fitness points in evolutionary phase 2, in environment shown in figure 1(b).

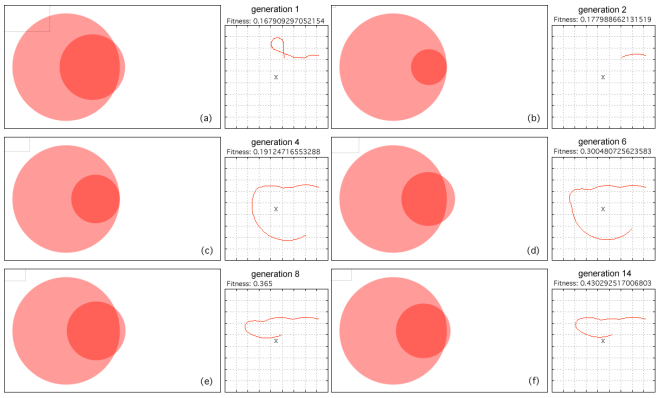

Figure 6: Area-proportional information relationship Venn diagrams and robot trajectories from individuals at key fitness points in evolutionary phase 3, in environment shown in figure 1(c).

walls (figure 1(b)). However, their behaviour and thus information profile was altered by this environmental shift, as illustrated in figure 5(a). The increased time spent navigating past the lengthened wall reduced overall repetition, and hence $H(X,Y)$ increased from the previous phase. This also shifted the proportion of information towards the sensors, and decreased the independence of motor state, reducing proportional $H(Y|X)$.

A major change occurred at generation 6 (figure 5(b)), with a complex network of gene interactions stopping the robot sharply shortly after it reaches the light. The robot then spends the rest of its lifetime entirely stationary, which heavily impacts on the system's information profile. The long sequence of repeated identical sensory and motor data events this generates sees $H(X,Y)$ drop sharply, $H(X)$ ends up close to $H(Y)$, and the proportion of mutual information between sensory and motor information is greatly increased.

Figure 6 shows key points from the third evolutionary phase. Individuals from the final generation of phase 2 perform poorly as a seed generation in this environment, getting stuck against the top of the broken wall, as shown in figure 6(a). Generation 2 sees the genetic network re-wired, almost entirely re-setting the system to a behaviour similar to that exhibited at the start of the first phase in figure 4(a), although with a GRN far more sophisticated in terms of inter-gene connectivity and use of sensory inputs, albeit producing dysfunctional behaviour at this point in time. This re-wiring also has a dramatic effect on the information profile. $H(X,Y)$ is tiny, as the system spends the vast majority of its time stationary, stuck against a wall. However, motor state is almost entirely determined by sensor state—the area $H(Y)$ appears inside $H(X)$, although this could be another feature of extended inactivity.

Generation 4, figure 6(c), breaks out of this behaviour, further re-wiring the already highly connected network to incorporate all available sensors. It navigates very rapidly past the wall, arcs past the light source and finally gets stuck against the side wall. The increase in activity brings an increase in joint entropy, although as with the rest of the individuals from this phase, a relatively long time spent stationary against a wall at the end of the run prevents $H(X,Y)$ reaching the levels seen during previous phases. The fact that $H(Y|X)$ is close to zero indicates that motor state is still almost entirely predictably mapped from sensor state, although the GRN is sufficiently complex that this mapping is difficult to verify.

Despite a noticeable increase in fitness, the GRN wiring remains unaltered in the transition to generation 6 (figure 6(d)), the only changes being a tiny shift in a single protein decay rate and a small change to the protein saturation level (maximum possible protein instances). This is a relatively common theme working with this GRN model: small changes to parameters that do not directly impact gene wiring can have a major impact on behaviour and thus fitness. The jump in fitness from 0.19 to 0.3 appears to be due the robot ending its run stuck at a more advantageous orientation.

There are then two more considerable gene connectivity changes in the shifts from generations 6 to 8 (figure 6(e)), and 8 to 14 (figure 6(f)), producing relatively minor behavioural shifts and correspondingly small changes to the information proportions.

Discussion

There appears to be an association between the degree of overlap of $H(X)$ and $H(Y)$, and the complexity of the environment in terms of the extent to which it presents an obstacle to the robot. This overlap may be due to the fact that as we move through phases 1–3, successfully achieving exposure to light depends increasingly on a tight relationship between sensor input and motor output. In the environment used in phase 3, figure 1(c), for example, the robot has to precisely and extensively alter its motor states based on fairly small fluctuations in proximity sensor data. Comparing all of the information proportional Venn diagrams across the three evolutionary phases, motor information indeed becomes increasingly determined by sensor information. Phase 2 may be skewed by the long period of inactivity exhibited by individuals however.

It would be interesting to increase the potential complexity of the GRN controllers (at least in terms of the number of genes, and possibly the number of proteins also) and environment to see if the trend is maintained. It may actually be that substantially more complex controllers in more complex task environments necessarily use *less* linear mappings from sensor data to actuator output. We predict that in such situations we would observe relatively low values for $I(X;Y)/H(X,Y)$ (proportion of the joint entropy that is mutual information), because for a given sensor input there will be uncertainty regarding the actuator output, for example if dependent on a temporally extended process of engagement between robot and environment. We would also expect relatively high values for $H(X,Y)$ again as a result of overall unpredictability. However, we would not expect to see $H(X,Y)$ approach its maximum possible value in any 'intelligent' system, as this equates to randomness.

Another fairly consistent if general relationship exists in the relative areas of $H(X)$ and $H(Y)$. It seems that in phase 3, again as a result of fit behaviour depending on motor states being tightly controlled by sensor data, the amount of information contained in the sensor data is much greater than that contained in the motor data. This may be enhanced because the motor data consists of two sets of wheel speeds—containing less potential for variation than the four sets of sensor data (light and proximity from the left and right sides of the robot) that make up $H(X)$. In phase 1 a similar relationship holds, possibly due to the simple and repetitive nature of the two wheel speeds when in a simple circular trajectory: contrast for example the less regular trajectory plots and slightly greater areas of figure 4(a) and (b) with (c)–(e).

We suspect that the relative areas of $H(X)$ and $H(Y)$ reflect sensor and actuator sophistication (in terms of potential information content), and the nature of the task environment, in that we would expect systems with similar degrees of sensor and actuator sophistication to show similar proportions for $H(X)$ and $H(Y)$, given a task environment that depends of extensive use of both. For example, and AIBO dancing to music should show a greater balance between the two than one merely tapping a paw, in which case the amount of sensor information would be expected to outweigh actuator information.

With regard to the results presented, it is difficult to tie variation in behaviour to variation in information proportion in a very precise way when this variation arises from controllers operating across identical 'bodies' and environments performing a general task consisting of moving about. We consider it unlikely that using real Khepera robots in the same configuration as used here would make much difference, just because the set-up is so simple that large discrepancies in sensor and motor data between the simulated and the real are unlikely.

A direction planned for future work is to perform information profiling across a range of different systems, as well as for more complex systems (more sensors, more Degrees of Freedom, etc.) in more complex environments (more interactive, greater plasticity, etc.), such as Sony AIBO robots operating in a physical environment.

Conclusions

We have illustrated the evolution of GRN-driven control systems for a simulated robot, demonstrating the ability of a simple 5-gene GRN to perform basic sensor integration and control. We have shown how information-theory based measures, particularly proportional measures of entropy, mutual information and conditional entropy can be used to analyse sensor and actuator data from such a system. We have suggested relationships between these measures and features of the systems under analysis, while recognising that additional work is necessary to prove any general principles. Our primary goal in the near future is to develop these measures further with a view to producing measures of system–environment information exchange as a basis for introducing quantification to earlier work characterising the notion of embodiment in just such relational terms (Quick et al., 1999).

Acknowledgments

The first author would like to thank Lars Olsson and Alex Klyubin for their wisdom and helpfulness on the subject of information theory.

References

Alberts, B. et al. (2002). *Molecular Biology of the Cell*. Garland, 4th edition.

Arthur, W. (2000). *The Origin of Animal Body Plans: A Study in Evolutionary Developmental Biology*. Cambridge.

Astor, J. C. and Adami, C. (2000). A developmental model for the evolution of artificial neural networks. *Artificial Life*, 6:189–218.

Bongard, J. C. (2002). Evolving modular genetic regulatory networks. In *Proc. IEEE 2002 Congress on Evo. Comp.*, pages 1872–1877. IEEE Press.

Chow, S. and Ruskey, F. (2004). Drawing area-proportional Venn and Euler diagrams. In Liotta, G., editor, *Graph Drawing, 11th Int. Sym., Revised Papers*, LNCS, pages 466–477. Springer.

Cover, T. M. and Thomas, J. A. (2001). *Elements of Information Theory*. Wiley, New York.

Davidson, E. H. (2001). *Genomic Regulatory Systems: Development and Evolution*. Academic Press.

de Jong, H. et al. (2003). Hybrid modeling and simulation of genetic regulatory networks: A qualitative approach. In Pnueli, A. and Maler, O., editors, *Hybrid Systems: Computation and Control*, volume 2623 of *LNCS*, pages 267–282. Springer-Verlag.

Dellaert, F. and Beer, R. (1996). A developmental model for the evolution of complete autonomous agents. In Maes, P. et al., editors, *From Animals to Animats 4: Proc. 4th Int. Conf. on Simulation of Adaptive Behaviour*, pages 393–401. MIT Press.

Eggenberger, P. (1997). Evolving morphologies of simulated 3D organisms based on differential gene expression. In Husbands, P. and Harvey, I., editors, *Proc. 4th European Conf. on Artificial Life*, pages 205–213. MIT Press.

Goldberg, D. (1989). *Genetic Algoriths in search, optimization and machine learning*. Addison-Wesley.

Hornby, G. S. and Pollack, J. (2002). Creating high-level components with a generative representation for body-brain evolution. *Artificial Life*, 8(3):223–246.

Kauffman, S. (1969). Metabolic stability and epigenesis in randomly constructed genetic nets. *Journal of Theoretical Biology*, 22:437–467.

Kumar, S. and Bentley, P. J. (2003). Biologically inspired evolutionary development. In *Proceedings of the 5th Internation Conference on Evolvable Systems (ICES): From Biology to Hardware*, LNAI. Springer.

Lungarella, M., Pegors, T., Bulwinkle, D., and Sporns, O. (2005). Methods for quantifying the informational structure of sensory and motor data. *Neuroinformatics*, 3:243–262.

Olsson, L., Nehaniv, C. L., and Polani, D. (2004). Sensory channel grouping and structure from uninterpreted sensor data. In *2004 NASA/DoD Conf. on Evo. Hardware*, pages 153–160. IEEE.

Quick, T., Dautenhahn, K., Nehaniv, C. L., and Roberts, G. (1999). On bots and bacteria: Ontology independent embodiment. In Floreano, D., editor, *Proc. 5th Euro. Conf. on Artificial Life*, LNAI 1674, pages 339–343. Springer.

Quick, T., Nehaniv, C. L., Dautenhahn, K., and Roberts, G. (2003). Evolving genetic regulatory network-driven control systems. In Banzhaf, W., editor, *Advances in Artificial Life. 7th Euro. Conf.*, LNAI 2801, pages 266–277. Springer.

Reil, T. (1999). Dynamics of gene expression in an artificial genome - implications for biological and artificial ontogeny. In Floreano, D., editor, *Proc. 5th Euro. Conf. on Artificial Life*, LNAI 1674, pages 457–466. Springer.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423, 623,656.

West-Eberhard, M. J. (2003). *Developmental plasticity and evolution*. Oxford University Press.