Bilaterally Symmetrical Encoding in the Evolution of Artificial Neural Networks for Symmetry Detection

Derek James
University of Louisiana at Lafayette
Lafayette, LA 70504
(337) 482-1130
djames@gmail.com

Philip Tucker
Dallas, TX 75206
(972) 297-3456
ptucker@gmail.com

Anthony Maida
University of Louisiana at Lafayette
Lafayette, LA 70504
(337) 482-6308
maida@cacs.louisiana.edu

ABSTRACT

In order for neuroevolutionary techniques to produce increasingly complex and sophisticated topologies, new methods need to be developed which effectively exploit reuse and modularity. Bilateral symmetry is an important form of reuse and a key feature of complex biological central nervous systems. We present a method for encoding bilateral symmetry within the context of an existing neuroevolutionary framework, NEAT (NeuroEvolution of Augmenting Topologies). We then present a model of symmetry detection that relies on the symmetry in the structure of the neural system to make symmetry judgments. We demonstrate that this model performs better than an asymmetrical representation on a symmetry discrimination task in which the axis of symmetry is given. On a second task, the networks must first find the axis of symmetry before making the symmetry judgment. In this task the symmetrical encoding performs worse than the asymmetrical one. We discuss some possible explanations for these results.

1.INTRODUCTION

The specific class of techniques that apply evolutionary algorithms to exploration of both the connection weights and topologies of artificial neural networks can be broadly classified in two ways. As Yao [18] described these classes of techniques: "At one extreme, all the details, i.e., every connection and node of an architecture, can be specified by the chromosome. This kind of representation scheme is called *direct encoding*." This type of encoding specifies a 1-to-1 mapping between a gene and a topological feature (such as a neuron or connection), which is both impractical from a scalability aspect and not biologically plausible.

Conversely, *indirect encoding* refers to the class of techniques in which there is some aspect of development used in deriving the network from the genetic representation so that there is not simply a direct 1-to-1 mapping of gene to feature. For example, a gene in an indirect encoding scheme may encode for a modular substructure that may be reused many times in the course of building the phenotype.

Stanley and Mikkulainen [17] refer to the development phase in such a computational paradigms as *artificial embryogeny* (AE). They point out the importance of genetic reuse and lay out the conceptual framework for classifying these different approaches. Modular reuse in neuroevolution has been investigated with a variety of different methods [1][8][12][14]. Eggenberger [5] explicitly selected for the emergence of bilateral symmetry in the evolution of artificial 3D morphologies and studied the effects. However, to our knowledge, the encoding of bilateral symmetry in artificial neural networks has not been investigated.

Bilateral symmetry is a key feature of complex biological nervous systems, and is realized very early in the development of body plans of many biological embryos. Since most of the sensory input for complex organisms is collected from bilaterally-symmetrical organs, and many motor tasks (such as locomotion) involve symmetrical appendages, it is reasonable to hypothesize that there is a correlation between the bilateral symmetry inherent in the organism's nervous system and the symmetry inherent in its sensory input and motor output.

It is also possible that the symmetrical structure of mammalian brains provides an explanation for other cognitive functions. In Braitenberg's *Vehicles* [2], he proposes a simple bilateral symmetry detector, an array of elements symmetrically spaced from each other with respect to a midline. An image projected onto such an array would exhibit strong activation for symmetrical images, in which input is balanced on both sides of the input array with respect to the midline, but weak activation for asymmetrical imagery, in which input is not balanced. Other researchers [4], including Braitenberg [3] have proposed that the symmetrical organization of the human brain may be responsible for the innate preference for vertical symmetry [5] and its increased salience over other orientations.

Palmer and Hemenway [13] proposed a two-stage model of symmetry detection in humans, the first stage involving the location of the axis of symmetry, and the second the actual symmetry judgment. In recent years, strong models of symmetry detection have been developed [7][9]. However, these cognitive models of symmetry detection do not draw a direct link between the symmetry inherent in the task and the underlying symmetry of the neural structure. Neural network models of symmetry detection [11][15] have found connection weights trained via backpropagation tend to converge on similar values for symmetrical connections, though they have been criticized for either being too small to scale realistically, or taking too many epochs to train. However, the authors of [10] assert that increased preference and performance in vertical symmetry is a result of more early exposure to stimuli with vertical axes of symmetry.

We introduce a new method of indirect encoding for the evolution of artificial neural networks that allows for the flexible exploration of topologies with varying degrees of bilateral symmetry. Using this encoding scheme, we introduce a model for symmetry detection with a basis on the underlying structure of neural systems.

GECCO'06, July 8–12, 2006, Seattle, WA, USA.
ACM 1-59593-186-4/06/0007.

2.BILATERAL SYMMETRY ENCODING

The new encoding system was developed as a modification to an existing neuroevolutionary algorithm, NEAT (NeuroEvolution of Augmenting Topologies) [16]. In its original implementation, NEAT is a direct encoding algorithm. There are only two types of genes specified: neuron genes and connection genes. Each gene encodes for exactly one topological feature, i.e., one neuron gene encodes for exactly one neuron in the resulting neural network.

In the new encoding, the chromosome provides all information necessary for deterministically constructing the network. These instructions are in the form of additional attributes for each gene that define sidedness and connectivity. Unlike previous implementations of NEAT, the resulting network has spatial structure: a midline, and right and left sides. Connections either cross the midline when connecting neurons (contralateral) or do not cross the midline (ipsilateral).

All genes are defined by a symmetry attribute with three values: LEFT, RIGHT, or BOTH. Neuron genes with the LEFT or RIGHT value are transcribed into neurons on that side of the resulting network. Neuron genes with the BOTH attribute are transcribed into mirrored copies on both sides of the resulting network (Fig. 1).

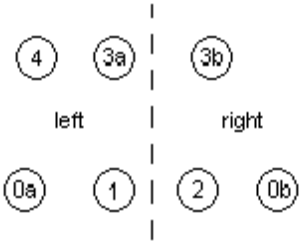

Figure 1. Neuron genes 1 and 4 have a symmetry attribute of LEFT and neuron gene 2 has a symmetry attribute of RIGHT. Neuron genes 0 and 3 have a symmetry attribute of BOTH, so they encode for mirrored copies (a and b).

In addition to the symmetry attribute, connection genes also have a *lateralization* attribute, with one of two values: IPSILATERAL (does not cross the midline) and CONTRALATERAL (crosses the midline). There are three cases of connectivity (where A denotes asymmetric neurons and S denotes symmetric neurons):

1) A to A
2) A to S and S to A
3) S to S

For A to A connectivity, symmetry and lateralization attributes are ignored and transcription proceeds as in previous implementations of NEAT (Figure 2).

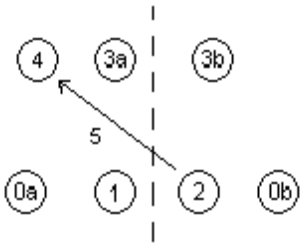

Figure 2. Connection gene 5 encodes for a connection between LEFT neuron 4 and RIGHT neuron.

Figure 3 shows the three possible cases of connectivity for A to S. In these cases, the symmetry attribute (LEFT, RIGHT, or BOTH) indicates which copies of the symmetrical neuron the connection will connect either to or from. The lateralization attribute is ignored.

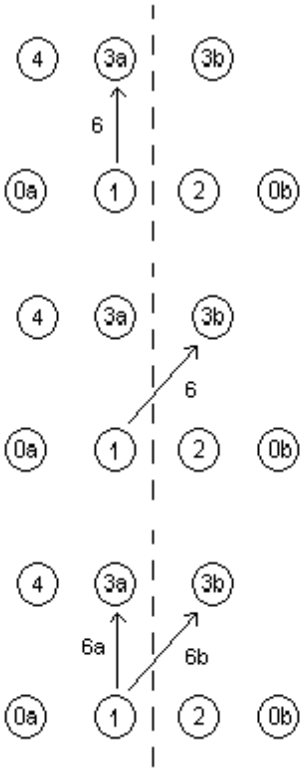

Figure 3. Three cases for connectivity between A to S neurons: left (top), right (middle), and both (bottom). These attributes encode for S to A connectivity in the same way.

Figure 4 shows the six possible cases of connectivity between pairs of symmetrical neurons. In these cases, the symmetry attribute defines the connectivity from the source neuron(s), so a connection gene with a LEFT value would only connect from the left copy of the source neuron. The lateralization attribute determines whether the transcribed connection(s) cross the midline or not.

For the initial topology, per the original NEAT paradigm, only input and output neurons are defined, which now may be either symmetrical or asymmetrical. NEAT also begins minimally, with no hidden units and only feedforward connections. Initially, all connections have a symmetry value of BOTH and a lateralization value of IPSILATERAL. This choice was made under the assumption that connectivity in biological brains is predominantly interhemispheric. This is one aspect of the implementation that is could be parameterized and explored in future work.

NEAT has three mutation operators: 1) Perturb existing connection weights, 2) Add new connection (including recurrent ones), and 3) Add new neuron. These mutations act on the genes, and work as before with the following exceptions.

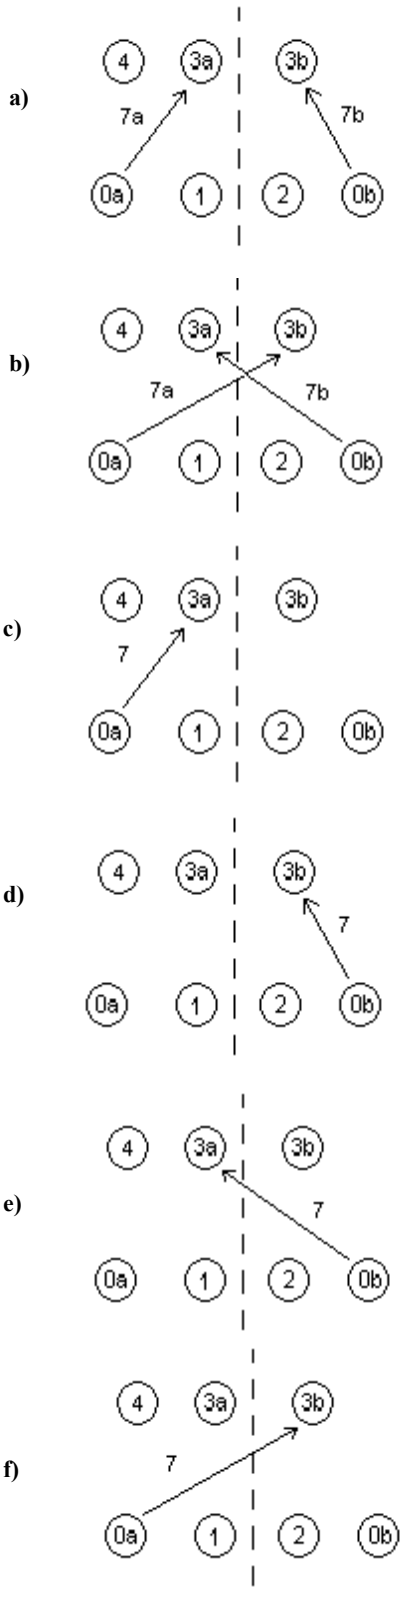

Figure 4. Connectivity between symmetrical neurons. (a) BOTH-IPSI, (b) BOTH-CONTRA, (c) LEFT-IPSI, (d) RIGHT-IPSI, (e) RIGHT-CONTRA, (f) LEFT-CONTRA.

In NEAT, new neurons mutate in the place of previous connections, actually creating three new topological features: the new neuron, a new incoming connection and a new outgoing connection. This works as before, and all three new features inherit the symmetry and lateralization attributes of the old connection.

If a new connection is added between A to S, S to A, or S to S neurons, a new symmetry rate determines the probability that the new connection will either be BOTH or LEFT/RIGHT. If the connection is between S to S neurons, another new symmetry rate determines the probability that the new connection is either CONTRALATERAL or IPSILATERAL. For all experiments described here, both of these rates were set at 0.50.

3.SYMMETRY DETECTION MODEL

As mentioned earlier, Braitenberg proposed a simple conceptual model for an idealized symmetry detector, as shown in Figure 5.

Figure 5. An array of simple elements upon which an image is projected. Elements symmetrical to each other with respect to the midline enhance each other. Adapted from [2].

The basic idea is that there is stronger activation of units if the projected image is symmetrical than if one side or the other is receiving more input. Contemporary models of bilateral symmetry detection do not explicitly refer to the structure of the underlying neural architecture.

The human visual system includes a much more complex, but similar architecture to Braitenberg's thought experiment. Visual input is projected onto regions in the back of the eye corresponding with a left and right visual field. This information is maintained in separate pathways, along the optic nerve, to symmetrical structures in each hemisphere of the visual cortex.

We propose a model of bilateral symmetry detection in humans in which input from two halves of a visual scene are input into mirrored halves of a symmetrical artificial neural network.

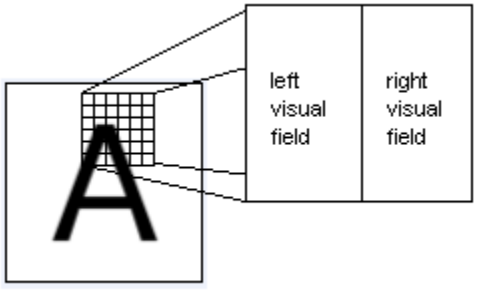

Figure 6. Pixel information is partitioned into left and right fields before being input into the network.

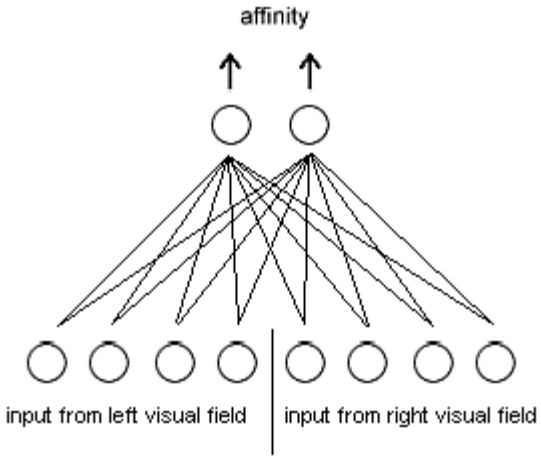

Figure 7. Network architecture for symmetrical encoding. Information from the image is input into corresponding symmetrical input neurons. Affinity is read from symmetrical output neurons.

Figure 7 shows the network architecture for the symmetrical encoding. Inputs from the corresponding visual fields are input into symmetrical input arrays. Symmetrical affinity outputs determine the symmetry judgment. Confidence was a function of the difference between the outputs. In the asymmetrical encoding, a single affinity output was used to make the confidence judgment.

Details of the network topologies for the two experiments will be elaborated on in the next section, but the basic model is one in which equal activation from balanced inputs from the visual stimuli will lead to balanced affinity outputs. The more similar the output values, the more likely a symmetry judgment will be made.

In one experiment, we incorporated movement into the behavior of the system, and the ability for the artificial retina to focus. However, the model does not account for the rotation of the axis of symmetry or for other types of symmetry.

4.EXPERIMENTAL SETUP

The two experiments were designed to test the performance of the two types of encoding (symmetrical and asymmetrical) under two conditions. Experiment 1 compares the performance of each when the axis of symmetry is already provided, which is comparable to Palmer and Hemenway's second stage of symmetry detection, the symmetry judgment. Experiment 2 compares the performance of each type of encoding for the two-stage model, in which the axis of symmetry must be located before the symmetry judgment is made.

Parameters for both experiments are shown in Table 1.

Experiment 1

In this experiment, the artificial retina was unable to move. The resolution of the retina was 6x6. Symmetrical and asymmetrical dot patterns were generated for use as stimuli. A sample set of the stimuli, which is very similar to that used in [10] is shown in Figure 8.

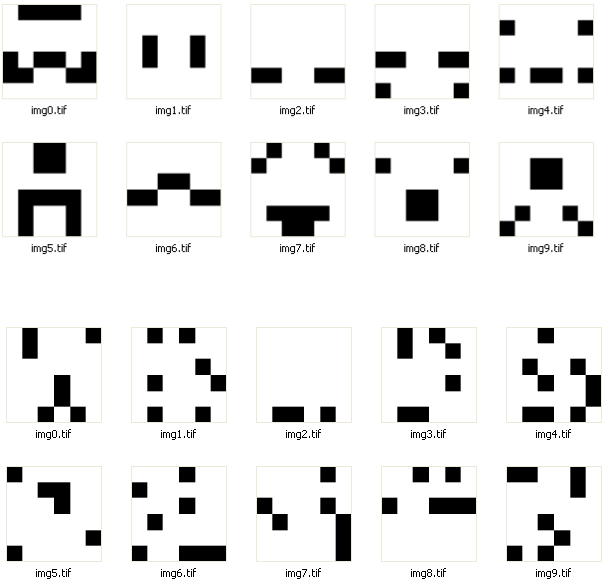

Figure 8. Sample of symmetrical (top two rows) and asymmetrical (bottom two rows) dot pattern stimuli used in Experiment 1.

The symmetrical patterns were generated by randomly toggling 10% of the pixels on one side of a white 6x6 canvas, then creating a mirror image. Asymmetrical stimuli were created by randomly toggling 20% of the pixels to black. 100 of each type of image were generated.

Every generation, each individual was viewed all 200 images, in random order, for 20 time steps. Affinity was calculated by a sum weighted toward outputs at later time steps, according to the formula in Figure 9.

$$\frac{\sum_{i=1}^{n} (\text{affinity}_i \cdot i^2)}{\sum_{i=1}^{n} i^2}$$

Figure 9. Equation for weighted affinity, were n = number of time steps, affinityi = affinity at step i.

If the weighted sum was greater than 0.50, the image was judged symmetrical. If less than 0.50, the image was judged asymmetrical. The fitness of each individual was the percentage of correct judgments.

There were two experimental groups, one using the bilaterally symmetric encoding described in previous sections, and an asymmetrically encoded group, using the original NEAT encoding. Ten runs were performed with each type of encoding.

Experiment 2

Symmetrical and asymmetrical letter stimuli were used for the second experiment. The full stimuli set can be seen in Figure 10.

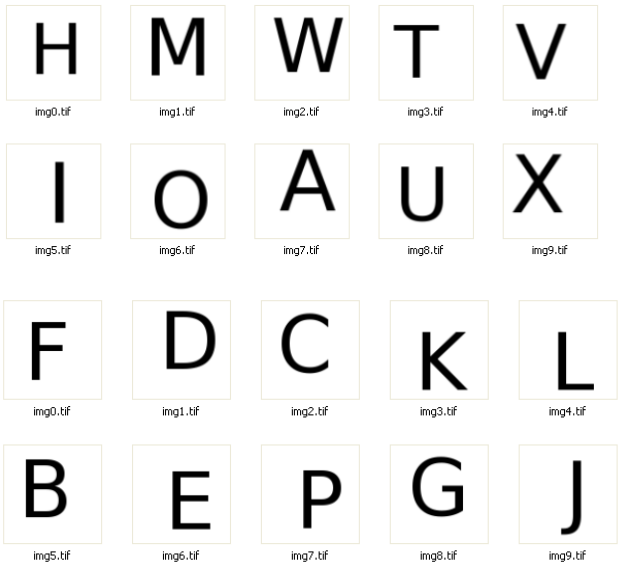

Figure 10. Samples of the 20 symmetrical (top two rows) and asymmetrical (bottom two rows) letter stimuli used in Experiment 2.

Each image was a 100x100 pixel grayscale TIFF with the letter centered on the canvas. For each individual, every trial, four randomized copies of each stimuli were presented, for a total of 40 symmetrical stimuli and 40 asymmetrical stimuli. The positions of the letters were randomly translated along the x and y axes, while ensuring that the entire letter remained visible on the canvas. Again, the order of presentation was randomized each trial.

In these experiments, each network had three additional output units, whose output determined a change in position along the x axis, y axis, and zoom factor respectively. Three additional inputs corresponding to the x, y, and z positions were also added to the networks. The artificial retina began each trial snapped to the border of the canvas, but each time step was allowed to actively scan the image. Each network was given 50 time steps to scan the image. Affinity judgments were calculated the same as in Experiment 1. Population size was 200 and survival rate was 20%. Parameters were identical to Experiment 1 except for the number of generations, which was 200.

5.RESULTS

The percentage correct for the fittest individual of the final generation was used to assess performance. Figure 11 shows the results for Experiment 1.

Populations with symmetrical encoding ($\underline{M}$ = 81.70, $\underline{SD}$ = 1.78) outperformed those with asymmetrical encoding ($\underline{M}$ = 77.85, $\underline{SD}$ = 1.49). This result was significant: $\underline{t}(18) = 5.24$, $\underline{p} < 0.05$.

The fittest individuals from the final generation of each run in Experiment 2 were evaluated on 100 randomizations of the symmetrical letter stimuli and 100 randomizations of the asymmetrical letter stimuli, with equal numbers of each letter. Figure 12 shows the results for Experiment 2.

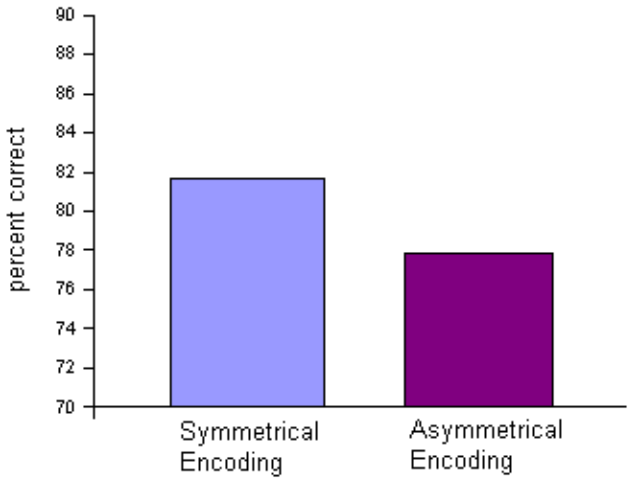

Figure 11. Results from Experiment 1. The mean of the percent correct for the fittest individual of the final generation for 10 runs is shown for each type of encoding.

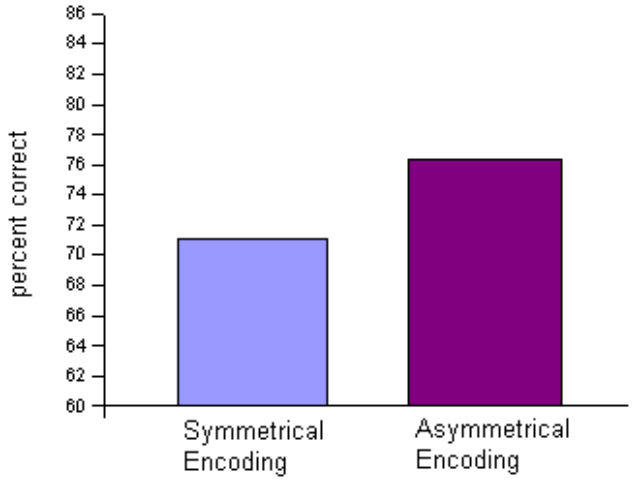

Figure 12. Results from Experiment 2. The percent correct for the fittest individual of the final generation on a test set of 100 random examples of each set of letter stimuli.

In this case, populations with asymmetrical encoding ($\underline{M}$ = 76.30, $\underline{SD}$ = 5.16) outperformed those with symmetrical encoding ($\underline{M}$ = 71.10, $\underline{SD}$ = 5.81). This result was significant ($\underline{t}(18) = 2.12$, $\underline{p} < 0.05$).

6.DISCUSSION AND FUTURE WORK

The results from Experiment 1 suggest that providing a means for symmetrical structure to easily evolve in the neural networks confers an advantage in symmetry judgment. In terms of Palmer and Hemenway's two-stage model, the evolved networks with symmetrical encoding perform better at the second stage, the symmetry judgment, than the asymmetrically encoded populations.

So why do the asymmetrically encoded populations do better in Experiment 2, which is analogous to the full two-stage process of

finding the axis of symmetry first, then making the symmetry judgment?

Some insight might come from an analysis of behavior of the fittest individuals from each of these populations. As a general trend, the asymmetrically encoded networks tended to zoom in on the image very early in each evaluation and focus on localized features (such as the intersection of two lines in the letter). By contrast, the symmetrically encoded networks tended to stay zoomed out through most of the trials. This suggests that the networks are attempting to solve the problem in different ways, that the asymmetrically encoded individuals are basing the discrimination on local features, while the symmetrically encoded individuals are basing their judgment on global features (such as symmetry).

We believe that one possible problem with Experiment 2 was the limited size of the stimulus set. There were enough differences based on local features that the networks could basically memorize the entire set, instead of making generalizations based on common global features. An obvious follow-up experiment would be to replicate Experiment 2 with a stimulus set similar in size to that used in Experiment 1.

However, we do feel confident that we have demonstrated that our symmetrical encoding does confer a benefit, not only in the efficiency and compactness of the genetic representations, but in performance of symmetry judgment.

In keeping with this line of research, it would be useful to test the encoding and symmetry model against other kinds of symmetry, such as those in which the axis is rotated, horizontal symmetry, and multiple symmetries. It would also be useful to test symmetrical encoding at other visual perception tasks, such as face recognition.

As for other domains, we would like to see how this type of encoding performs in the context of goal-directed agent control. Most mobile agents feature symmetrical sensory input as well as symmetrical body plans. This seems like a logical domain in which to test whether symmetrical encoding demonstrates further benefits.

Finally, while symmetrical encoding is an important type of indirect encoding, it is one of many early steps. We would like to explore other types of indirect encoding, modular encoding in particular, and build a developmental toolbox for evolving increasingly complex neural structures in order to solve more demanding real-world problems and provide further insight into biological cognition.

7.REFERENCES

[1] Bongard, J.C. Evolving modular genetic regulatory networks. *Proceedings of Congress on Evolutionary Computation* (2002), 1872-1877.

[2] Braitenberg, V. *Vehicles*. The MIT Press, Cambridge, Massachussetts, 1984.

[3] Braitenberg, V. Reading The Structure of Brains, *Network*, (1989), 1:1-11.

[4] Corballis, M. C. and Roldan, C. E. Detection of symmetry as a function of angular orientation. *Journal of Experimental Psychology: Human Perceptual Performance*. (1975), 1: 221-230.

[5] Eggenberger, P. Evolving morphologies of simulated 3d organisms based on differential gene expression. *Proceedings of the 4th European Conference on Artificial Life*, (1997), 205-213.

[6] Evans, C. S., Wenderothô, P., Cheng, K. Detection of bilateral symmetry in complex biological images. *Perception*, (2000), 29: 31-42.

[7] Ferguson, R. W. MAGI: Analogy-based encoding using symmetry and regularity. *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, (1994), 283-288.

[8] Gruau, F. Genetic Synthesis of Modular Neural Networks. *Proceedings of the Fifth International Conference on Genetic Algorithms*, (1993), 318.

[9] Labonte, F., Shapira, Y., Cohen, P., and Faubert, J. A model for global symmetry detection in dense images. *Spatial Vision* (1995), 9: 1, 33-55.

[10] Latimer, C., Joung, W., and Stevens, C. Modelling symmetry detection with back-propagation networks. *Spatial Vision* (1994), 8:4, 415-431.

[11] McClelland, J. L., Rumelhardt, D. E. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. (1986) MIT Press, Cambridge, MA.

[12] Nolfi, S. Using emergent modularity to develop control systems for mobile robots. *Adaptive Behavior*, (1997), 5: 343-363.

[13] Palmer, S.E. And Hemenway, K. Orientation and symmetry: Effects of multiple, rotational, and near symmetries. *Journal of Experimental Psychology: Human Perception and Performance*, (1978), 16: 150-163.

[14] Reisinger, J., Stanley, K.O. and Miikkulainen, R. Evolving Reusable Neural Modules, *Proceedings of the Genetic and Evolutionary Computation Conference* (2004). New York, NY: Springer-Verlag.

[15] Sejnowski, T. J., Kienker, P. K., and Hinton, G. E. Learning symmetry groups with hidden units: Beyond the perceptron. *Physica* (1986), 22D: 260-275.

[16] Stanley, K.O. and Miikkulainen, R. Efficient Evolution of Neural Network Topologies, *Proceedings of the 2002 Congress on Evolutionary Computation*.

[17] Stanley, K. O. and Miikkulainen, R. A Taxonomy for Artificial Embryogeny. *Artificial Life*, (2003) 9(2): 93-138.

[18] Yao, X. Evolving artificial neural networks, *Proceedings of the IEEE*, 87(9):1(1999), 423-1447.