On Preferred States of Agents: how Global Structure is reflected in Local Structure

Tom Anthony[1], Daniel Polani[1,2] and Chrystopher L. Nehaniv[1,2]

[1]Adaptive Systems Research Group
[2]Algorithms Research Group
School of Computer Science, University of Hertfordshire
College Lane, Hatfield, Herts, AL10 9AB, UK
{T.C.Anthony,D.Polani,C.L.Nehaniv}@herts.ac.uk

Abstract

We investigate the correlation between the information theoretic measure of *empowerment* and the graph theoretic measure of *closeness centrality*, to better understand the structural conditions that must exist in a world for learning and adaptation. We examine both measures in both a simple grid-world scenario, represented as a graph, and on a scale-free graph. We show a strong correlation between the two measures, and discuss the strengths and weaknesses of both. We go on to show how the local measurement of empowerment can in many cases predict a measure for the global measurement of closeness centrality.

Motivation

"Nature uses only the longest threads to weave her patterns, so that each small piece of her fabric reveals the organization of the entire tapestry." - Richard Feynman

Learning and adaptation are central themes to artificial life, and it is our hypothesis that a better understanding of the conditions that must exist to make learning, adaptation and evolution possible will help to guide future research. It is plausible to assume that an arbitrary or random world would be extremely difficult, if at all possible, to learn. We know there is significant structure in the world, and believe that learning takes advantage of this structure. In this paper we begin to investigate what conditions, embedded within a world through some underlying structure, are necessary for certain types of adaptation problems.

It has been hypothesised that embodied agents receive an adaptive and evolutionary advantage by optimising their sensoric and neural configurations for their environment. Specific attention has been paid to processing and optimising of Shannon-type information they receive from their environment (Attneave, 1954; Barlow, 1959, 2001; Atick, 1992). Similar work includes the concept of *homeokinesis*, proposed by Der et al. (1999), where a homeokinetic system, or agent, learns to improve the predictive capabilities of its future perceptions.

A specific flavour of this view suggests that such informational predictive principles could provide organisms/agents with intrinsic motivation. Examples include that by Prokopenko et al. (2006), Bialek et al. (2001) and Ay et al. (2008), which use similar approaches based on *excess entropy / predictive information*.

In this paper we have chosen to use *empowerment* (Klyubin et al., 2005b,a), an information theoretic measure for the efficiency of a *perception-action loop*. Essentially empowerment uses the channel capacity for the external component of a perception-action loop to identify areas that are advantageous for an agent embodied within an environment.

It assumes situations with a high efficiency of the perception-action loop should be favoured by an agent. Based entirely on the sensors and actuators of an agent, empowerment intrinsically encapsulates an evolutionary perspective; namely that evolution has selected which sensors and actuators a successful agent should have, which in turn suggests which states should be visited.

This hypothesis was tested in a variety of different scenarios (Klyubin et al., 2005b,a; Capdepuy et al., 2007), and notwithstanding the quite different scenarios it coincided surprisingly well with an intuitive understanding of favourable behaviours or of natural solutions to particular challenges of adaptation. Furthermore, it correlated well against measures that had been hand crafted to evaluate certain scenarios.

Notwithstanding the successful performance, we do not currently have a strong understanding of why this may be. What are the properties of the world that make empowerment such a universal measure? Why should it work at all? These are the questions we are going to begin to study in this paper.

Locating Structure

We hypothesise that an agent that optimises its sensorimotor apparatus improves its ability to detect the underlying structure of the world, and that this is an important aspect of such optimisation. We further hypothesise that a better understanding of this structure would improve such optimisation, and thus allow for better adaptation and learning.

To investigate this we set out to start identifying the basic properties of the world, and how they are detected by empowerment. We selected to go about this by investigating a representation for an environment that manifests its structure in an easily observable manner, is well understood, and has established methods for measuring preferable states.

We chose to represent the state space using graphs, which fit all these criteria; they are well understood through graph theory and social network analysis, and they have accessible methods for identifying certain aspects of their structure. As a measure to identify preferred states we chose to use *centrality*, a measure of a node's importance from graph theory, which is a well established method (Wasserman and Faust, 1994). There are varying measures for centrality; in this paper we use *closeness centrality*, which most closely corresponds with the spirit of empowerment.

Most stationary worlds, containing an embodied agent, can be viewed of as the current state of the world connected to neighbouring states by the actions the agent would need to take to arrive at them; this can be modelled as a graph. This same representation of the world was used by Şimşek and Barto (2007) in investigating skill development among agents.

We can now analyse empowerment, and some aspects of what it captures about the world, by comparing it with centrality measurements in the same scenarios.

Quantifying Preference

Empowerment, a local measure, quantifies the changes that an embodied agent can make on its environment, and observe the effects of, in a given time period. Here we reduced ourselves to a simple representation of the world which is entirely deterministic, creating a special case for empowerment. However, it can work in both entirely deterministic and probabilistic environments, which may even be non-stationary (Capdepuy et al., 2007).

The closeness centrality of a node in a graph is calculated by adding the distance of the shortest paths from that node to every other node in the network, and then inverting this value so that a shorter total path to all other nodes has a higher value. To calculate the closeness centrality of a signal node requires viewing the whole graph; it is a global measure. Klyubin et al. (2005a) showed an example where a similar measure, the average shortest distance in a maze, correlated well with empowerment.

We will examine two scenarios, and will employ both empowerment and closeness centrality in each for identifying and measuring states that an embodied agent would find 'interesting' or 'preferential' to be in. When we use the word 'state' we refer to the state of the whole system, including both the environment and the agent.

Information Theory

The notion of empowerment is based on information theory, introduced by Shannon (1948). To introduce this, the first

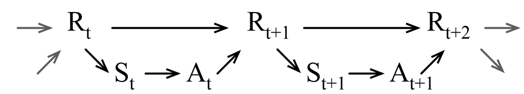

Figure 1: Bayesian network representation of the perception-action loop.

important measure is *entropy*, which is a measure of uncertainty:

$$H(X) = -\sum p(x) \log p(x). \tag{1}$$

Where X is a discrete random variable with values $x \in X$ and $p(x)$ is the probability mass function such that $p(x) = Pr\{X = x\}$. The logarithm can be taken to any chosen base; in our paper we consistently use 2, and accordingly the units of measurement are then called *bits*. If Y is another random variable jointly distributed with X the *conditional entropy* is:

$$H(Y|X) = -\sum_x p(x) \sum_y p(y|x) \log p(y|x). \tag{2}$$

This measures the remaining uncertainty about the value of Y, if we know the value of X. Finally, this also allows us to measure the *mutual information* between to random variables:

$$I(X;Y) = H(Y) - H(Y|X). \tag{3}$$

Mutual information can be thought of as the reduction in uncertainty about the variable X or Y, given that we know the value of the other. The mutual information is symmetric, so we could also use $I(X;Y) = H(X) - H(X|Y)$ (Cover and Thomas, 1991).

Empowerment

Empowerment is based on the information theoretic perception-action loop formalism introduced by Klyubin et al. (2005a, 2004), as a way to model embodied agents and their environments. The model views the world as a communication channel; when the agent performs an action, it is injecting Shannon information into the environment, which may or may not be modified, and subsequently the agent reacquires part of this information from the environment via its sensors.

In Fig.1 we can see the perception-action loop represented by a Bayesian network, where the random variable R_t represents the state of the environment, S_t the state of the sensors, and A_t the actuation selected by the agent at time t. It can be seen that R_{t+1} depends only on the state of the environment at time t, and the action just carried out by the agent.

By modelling this as a communication channel, we can employ information-theoretic methods, which are the basis for empowerment. First, we must introduce channel capacity (Shannon, 1948; Cover and Thomas, 1991) for a discrete memoryless channel:

$$C(p(y|x)) = \max_{p(x)} I(X;Y). \tag{4}$$

The random variable X represents the distribution of messages being sent over the channel, and Y the distribution of received signals. Clearly, the higher the mutual information between the two variables, the higher the capacity of the channel. The channel capacity is measured as the maximum mutual information taken over all possible input distributions, $p(x)$, and depends only on $p(y|x)$, which is fixed. One algorithm that can be used to find this maximum is the iterative Blahut-Arimoto algorithm (Blahut, 1972).

Empowerment can be intuitively thought of as a measure of how many observable adjustments an embodied agent can make to his environment, either immediately, or in the case of n-step empowerment, over a given period of time. An alternative way to view empowerment is that it guides agents to places in the world where they get the most benefit from their sensors and actuators. Using the above perception-action loop formalism and the Blahut-Arimoto algorithm, this can be directly quantified. We remind the reader that sensors and actuators implicitly encode evolutionary knowledge of the type of information to perceive and 'create'.

In the case of n-step empowerment, we first construct a compound random variable of the last n actuations, labelled A_t^n. We now need to maximise the mutual information between this variable and the sensor readings at time $t+n$, represented by S_{t+n}. Here we consider empowerment as the channel capacity between these:

$$\mathfrak{E} = C(p(s_{t+n}|a_t^n)) = \max_{p(a_t^n)} I(A_t^n; S_{t+n}). \tag{5}$$

An agent that maximises its empowerment will position itself in the environment in a way as to maximise its options for influencing its relationship with the environment (Klyubin et al., 2005a).

Note that in this paper we are use empowerment in an exclusively deterministic scenario, within a discrete world, but that empowerment is defined in full generality for non-deterministic probabilistic environments and does not assume perfect information.

In this paper we can use a shorthand method for calculating empowerment; we are able to do this for several reasons. All the scenarios we examine are deterministic and feature no non-stationary elements, and so do not require the probabilistic elements of empowerment. Additionally, as they are all represented as a graph, we are able to further simplify the formula. We can calculate n-step empowerment for a node v_i on the graph thus:

$$\mathfrak{E}_n(v_i) = \log \left[\sum_{\substack{j=1 \\ d(v_i,v_j)\leq n}}^{g} 1 \right] \tag{6}$$

Where $d(v_i, v_j)$ is the geodesic distance between the nodes v_i and v_j. Note that this is a shorthand method we are able to use as we have complete knowledge of the scenarios and the representation; Eq. (5) reduces to Eq. (6), and would work identically in the same scenarios, using the perception-action loop formalism.

Closeness Centrality

Graph Theory and Network Analysis have long had a requirement for identifying important nodes in a graph (Wasserman and Faust, 1994). The simplest methods for this have been to count the edges leaving or entering a node, known as outdegree and indegree respectively. This is very simplistic and is normally inadequate for complex graphs. Therefore, the primary method for measuring node importance is a group of various measures collectively known as centrality. There have been several methods of centrality suggested over time, but one of the most popular is closeness centrality, which can be presented in various ways. As mentioned in Wasserman and Faust (1994), and reviewed by Freeman (1979), the simplest formula for closeness centrality is that suggested by Sabidussi (1966):

$$C_C(v_i) = \left[\sum_{\substack{j=1 \\ j\neq i}}^{g} d(v_i, v_j) \right]^{-1}. \tag{7}$$

For a given node v_i, in a graph with g nodes, this gives a measurement of the sum of the shortest paths to all other nodes, which is then inverted to give a higher centrality to those with shorter total paths to the rest of the graph. Intuitively, this can be closely linked to the average distance from all other cells that empowerment was anti-correlated with, from the maze scenario used in Klyubin et al. (2005a).

To calculate the closeness centrality on the graphs encountered throughout this paper, we used the network analysis software Pajek (Batagelj and Mrvar, 1998). Pajek uses a modified version of closeness centrality, suggested in Beauchamp (1965):

$$C'_C(v_i) = \frac{(g-1)}{\left[\sum_{j=1}^{g} d(v_i, v_j)\right]} = (g-1)C_C(v_i). \tag{8}$$

This formula is used simply to normalise the closeness centrality figures to the graphs size in order to allow comparison of the figures between graphs of different sizes.

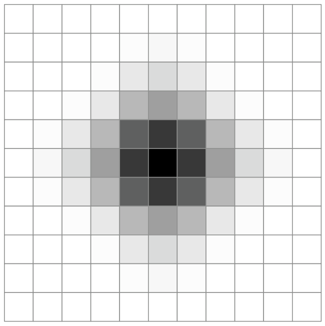

Figure 2: View of the empowerment distribution for the gridworld scenario, with the box positioned at the center. A darker shade means higher empowerment. Empowerment scales from 5.92 to 7.79 bits.

Scenarios

In order to compare these two measurements we apply them to the same two agent scenarios to identify the correlation between them, and any areas of disparity. In order to construct the first scenario, it is necessary to observe that most state spaces encompassing an agent in a stationary world can be naturally represented as a graph of nodes, with transitions leading between them corresponding to the actions of an agent within that world.

Box Pushing

Consider the box pushing scenario from Klyubin et al. (2005a) as a graph. The scenario consists of a gridworld of infinite size, within which there exists an agent and a box, each of which occupy a single cell. The box is visible to the agent; his view of the world consists of his position and the position of the box. The agent has 5 actions available to it at any time; it can stand still, or move to one of the four neighbouring cells. If the agent moves into a cell that is occupied by the box then the box is pushed, in the same direction, into the adjacent cell.

In Klyubin et al. (2005a) it was shown that for any n-step empowerment, the agent prefers being near the box, which gives it more influence on the state of the world. It most 'enjoyed' beginning on top of the box, where moving in and of the 4 directions would allow it to fall down next to the box, from where it could start pushing it like normal; this could be used as a starting position but was a position impossible for it to return to.

In translating this world into a graph representation, we needed to limit our originally infinite world to a finite graph. We investigate the influence of this finiteness by examining the growth of centrality. We show that beyond a certain horizon it can be seen that the centrality increases in a continuous fashion and that the centrality for the nodes represented in previous approximations grows proportionately. Whilst we do not offer a proof of this fact, in Fig.3 we demonstrate the point by showing the correlation between graph representations of increasing diameters.

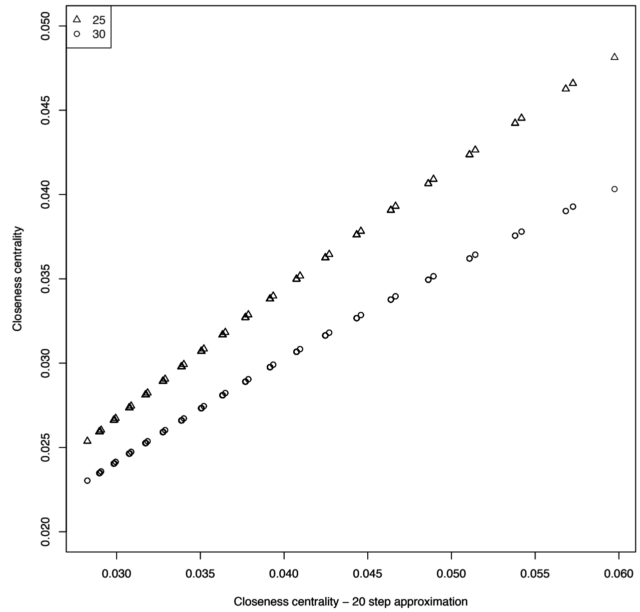

Figure 3: Correlation of closeness centrality for 25-step and 30-step graph approximations against a 20-step approximation.

Results

Klyubin et al. (2005a) had previously shown how empowerment worked in the box pushing gridworld experiment, and so it made for a good environment in which to run our initial experiments. We generated a unweighted directed graph to represent the world. Note that we are using a non-classical view of graphs; rather than viewing them as comprised of units, with connecting links between them, we are viewing each node as a possible state of the world, including the agent itself, (of which, only one can be the real state at any moment) and the edges as transitions between these states.

To do this, we initialised the world with the box in the center, and the agent standing upon the box, as described earlier. We then let the agent run through every possible trajectory of 30 actuations, generating a graph of states and actions; the final graph had 419,121 nodes. Using Pajek, we calculated the closeness centrality for all nodes in the graph.

We next measured empowerment for every state with the box positioned in the center of the world, and the agent positioned at each location that it could reach within 30 timesteps from the center. This was sufficient as the dynamics of the world comes from the agent's initial position relative to the box, and thus moving the box was unnecessary. Our empowerment measurements were run to measure 3-step, 5-step and 7-step empowerment.

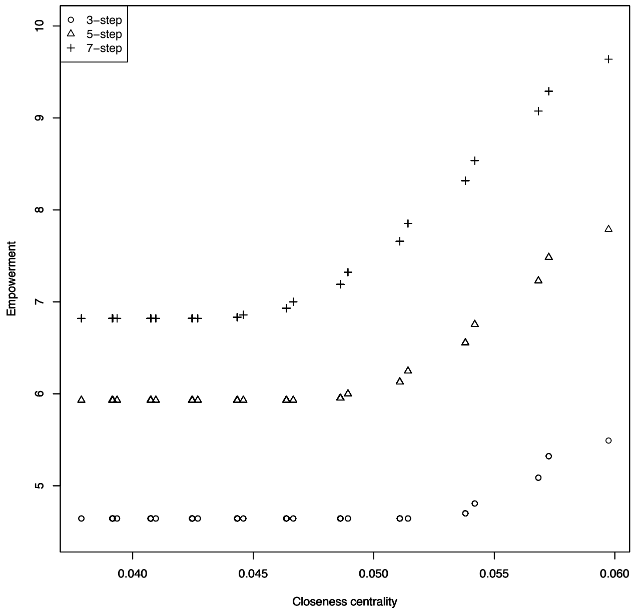

Figure 4: Correlation plot between Empowerment and Closeness Centrality. The horizon effect of empowerment can be seen clearly.

In order to correlate empowerment and centrality, we collated the results, removing the centrality results for nodes where the box was not positioned in the center of the world; this gave us a state for state comparison of each measure against the other for different initial positions of the agent.

We additionally ran the same experiment for graphs produced for both 20 and 25 timesteps, to identify the influence of representing the infinite gridworld as a finite graph did not skew the results. We found that the correlation of centrality for the overlapping nodes of these varying size graphs indicates a close to linear relationship and finite graphs work as a good approximation.

Note that closeness centrality is a global property, calculated it for any given node requires seeing all other nodes in the graph, while empowerment is local and looks only at neighbouring nodes within a given distance.

Local Structure

As hypothesised, we found a very strong correlation between the closeness centrality and empowerment, which can be seen in Fig.4. The graph shows clearly the horizon effect of empowerment; it can be seen to be constant whenever the box is outside of the agent's reach. For n-step value with larger values of n the horizon can be seen to extend further from the box. Once the box is within it's reach, according to n, the empowerment grows as the agent increases its influence over the world by getting closer to the box.

The horizon effect emphasises that empowerment is a local measure; it cannot see the whole world. However, when the agent is within an area where it can improve it's ability to manipulate the state of the world, this local measure correlates with the global measure of the world given by closeness centrality.

This highlights that in an infinite, or an unexplored, world where centrality cannot be employed, empowerment provides a measure that can be used. Whilst empowerment is limited by the horizon effect, exploring the world (which would be necessary to use closeness centrality) would allow our agent to also overcome the horizon.

In addition, this correlation also confirms our hypothesis that empowerment, within its horizon, does see global aspects of a system at a local level within this world. What structure or prerequisites that must exist for this effect to take place are yet to be determined.

It is important to note that the results from empowerment can be computed by the formula in Eq. (6), or equally by that in Eq. (5), without modelling the world as a graph at all.

Scale-free Graphs

The second scenario uses scale-free networks (graphs); a very important subclass of graphs, in which there are a few nodes with a high degree, and most nodes have a far lower degree. Their typical structure is independent of the graph's size; with fewer or more nodes, the graph would still exhibit similar properties. The exact distribution of edges per node follows a power law distribution (Barabasi and Albert, 1999):

$$P(k) \sim k^{-\gamma}. \tag{9}$$

Here $P(k)$ is the probability that a node connects with k other nodes, and decreases exponentially according to the coefficient γ.

As discussed in Barabasi (2003), scale-free graphs can be seen in many real world situations, including protein interaction networks (Jeong et al., 2001), social networks, and even the world wide web (Barabasi and Albert, 1999).

We hypothesise that the scale-free property of graphs can work to synthesise an underlying structure that may be found in real world task spaces, and can be used as a good platform for initial investigation of such structure.

Results

Using preferential attachment algorithm introduced by Barabasi and Albert (1999) we constructed a scale-free undirected graph with 400,000 nodes to run our measures on. Our graph was built using an initial complete graph of 3 nodes, and adding additional nodes one at a time. Each new node would create 3 new edges connected to 3 different nodes on the existing graph, chosen using a probability according to their current degree.

For all nodes in the graph we calculate both the n-step empowerment (for a range of values of n) and the closeness

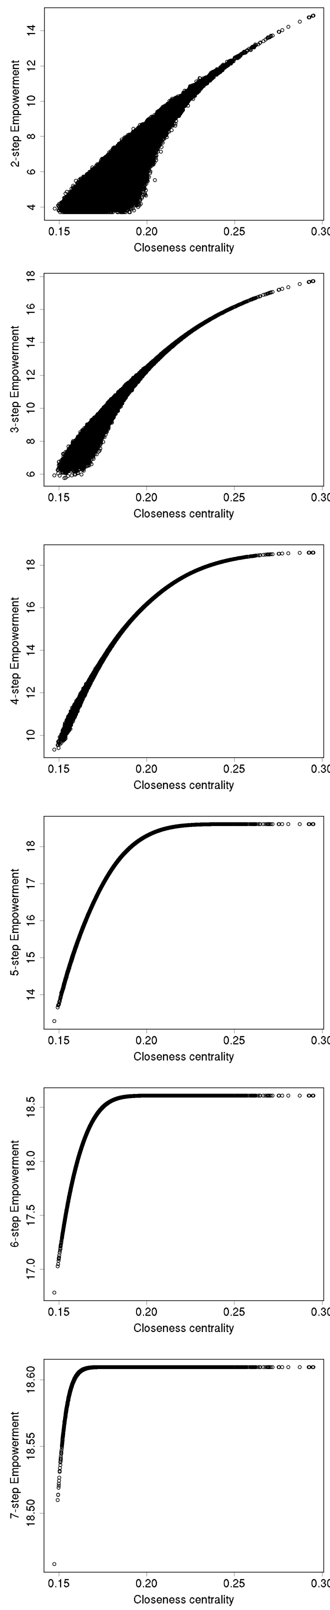

Figure 5: Correlation between closeness centrality and 2-step to 7-step empowerment.

centrality. To calculate the closeness centrality, we again use the Pajek analysis software.

Our results here corroborate those from our first experiment with regard to the correlation between closeness centrality and empowerment. Here, we see the inverse of the horizon effect; given too much time, empowerment can reach any part of the graph (analogous to being able to do anything within a world) and assigns almost all nodes equal value. This is an interesting point for empowerment; given too high of a 'budget', where an agent can do everything possible within the world (or reach every node in a graph) then it does not differentiate between them. This is the type of world we would describe as 'boring'; one where and agent can do anything it wants from any position of the world.

Again though, empowerment sees at a local level aspects of the global property of the world. In this scenario, this is maybe not surprising given the nature of a scale-free graph; but it is important to see that empowerment was not told anything of the structure of the world, and that still this fact comes through.

In Fig.5 we show the correlation between closeness centrality and n-step empowerment for n=2 to n=7. Note that even 2-step empowerment has a strong correlation at the higher centrality nodes, and 3-step even more so. As n increases it can be seen that the small-world property of the graph results in an empowerment ceiling being reached which results in a reduced correlation for high centrality nodes.

Discussion

Both of our experiments highlight the strong correlation between empowerment and closeness centrality, and that even n-step empowerment with a low value for n will normally serve a a strong predictor for centrality. This is significant given that individual node centrality is a global property of a graph, but we can use a local measure to give similar relative values to nodes. Note that empowerment doesn't see any more than centrality, but in the 'interesting' parts of the world it does see, the two measures agree.

In both scenarios the correlation is strong provided that the n chosen for n-step empowerment is suitable. We believe a simple method for overcoming this in an unknown world is for an agent to select the lowest n value possible; if the horizon of this n does not allow the agent to observe any degrees of freedom it can then increase n incrementally to overcome this (or embark on a random exploration).

With empowerment, selection of a suitable n is interesting in another regard; a low value of n can mean encountering the horizon effect, and possibly not seeing 'interesting' parts of the world, whilst a too high value of n can result in the agent being able to do anything and not needing to distinguish between different states. The result of this is a particular world having an n value with the correct balance between these two effects, which we hypothesise may reflect one aspect of the underlying structure which is important for learning and adaptation.

Closeness centrality is limited to deterministic task spaces that can be completely represented by either a directed or undirected graph, which constricts the space of problems it can be used to measure. In the space of problems in which both measures can be used, these results indicate not only empowerment correlates well with centrality, but it does so without complete knowledge of the world. Furthermore, it can work in non-deterministic, non-stationary, environments which cannot be represented as a graph, including infinite worlds.

A comparison could have been drawn between the global measure of closeness centrality, and some local version of centrality that worked on a local subset of the graph, and we expect a similar correlation would have been found. However, any such localised version of centrality would suffer from many of the same restrictions that centrality does compared to empowerment. We studied empowerment specifically as part of a much more general picture which includes an evolutionary aspect and which in addition will allow us to extend the research into non-deterministic environments in future work. Essentially we are using centrality as a 'sanity check' that empowerment does something sensible in these scenarios.

Overall, we believe that these results show a strong indication of certain global aspects of various worlds being 'coded' at a local level, and an appropriate sensory configuration can not only detect this information, but can also use it. Such uses could include learning and adaptation, and uses for evolution between generations. There are indications that understanding which aspects of global structure are visible at a local level would allow improved adaptation and learning for agents embodied within the corresponding world.

Future work

Vergassola et al. (2007) drew a parallel between the behaviour of biological organisms and search methods that use local informational cues to draw conclusions about the global structure of the world. It is our belief that further study of this area will allow us to not only draw further parallels with the learning and adaptation methods employed by biological organisms, but will also allow a better understanding of these processes leading to improved methods.

Further work needs to be done to extend these results into other worlds and task spaces, and to better understand in which scenarios they hold true. This should include worlds with various elements providing opportunities for agents to manipulate their environment, and even non-stationary worlds.

Attention needs to be paid to how to choose an initial strategy when presented with a completely unknown task space (such as choosing an initial n for empowerment) and conversely, how much of this information is embedded with an agent or organisms embodiment.

Acknowledgements

The authors would like to thank the reviewers for their feedback and suggestions which helped bring about various improvements to the paper.

References

Atick, J. J. (1992). Could information theory provide an ecological theory of sensory processing. *Network: Computation in Neural Systems*, 3(2):213–251.

Attneave, F. (1954). Some informational aspects of visual perception. *Psychological Review*, 61(3):183–193.

Ay, N., Bertschinger, N., R. Der, F. G., and Olbrich, E. (2008). Predictive information and explorative behavior of autonomous robots. *European Physical Journal B*. (Accepted).

Barabasi, A.-L. (2003). *Linked: How Everything Is Connected to Everything Else and What It Means for Business, Science, and Everyday Life*. Plume Books.

Barabasi, A. L. and Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439):509–512.

Barlow, H. B. (1959). Possible principles underlying the transformations of sensory messages. In Rosenblith, W. A., editor, *Sensory Communication: Contributions to the Symposium on Principles of Sensory Communication*, pages 217–234. The M.I.T. Press.

Barlow, H. B. (2001). Redundancy reduction revisited. *Network: Computation in Neural Systems*, 12(3):241–253.

Batagelj, V. and Mrvar, A. (1998). Pajek – program for large network analysis.

Beauchamp, M. A. (1965). An improved index of centrality. *Behavioral Science*, 2:161–163.

Bialek, W., Nemenman, I., and Tishby, N. (2001). Predictability, complexity, and learning. *Neural Comp.*, 13(11):2409–2463.

Blahut, R. (1972). Computation of channel capacity and rate distortion functions. *IEEE Transactions on Information Theory*, 18(4):460–473.

Capdepuy, P., Polani, D., and Nehaniv, C. L. (2007). Maximization of potential information flow as a universal utility for collective behaviour. In *Proceedings of the First IEEE Symposium on Artificial Life*.

Cover, T. M. and Thomas, J. A. (1991). *Elements of information theory*. Wiley-Interscience, New York, NY, USA.

Der, R., Steinmetz, U., and Pasemann, F. (1999). Homeokinesis - a new principle to back up evolution with learning. In Mohammadian, M., editor, *Computational Intelligence for Modelling, Control, and Automation*, volume 55 of *Concurrent Systems Engineering Series*, pages 43–47. IOS Press.

Freeman, L. C. (1979). Centrality in social networks: Conceptual clarification. *Social Networks*, 1(3):215–239.

Jeong, H., Mason, S. P., Barabasi, A. L., and Oltvai, Z. N. (2001). Lethality and centrality in protein networks. *Nature*, 411(6833):41–42.

Klyubin, A. S., Polani, D., and Nehaniv, C. L. (2004). Organization of the information flow in the perception-action loop of evolved agents. In Zebulum, R. S., Gwaltney, D., Hornby, G., Keymeulen, D., Lohn, J., and Stoica, A., editors, *Proceedings of 2004 NASA/DoD Conference on Evolvable Hardware*, pages 177–180. IEEE Computer Society.

Klyubin, A. S., Polani, D., and Nehaniv, C. L. (2005a). All else being equal be empowered. In Capcarrère, M. S., Freitas, A. A., Bentley, P. J., Johnson, C. G., and Timmis, J., editors, *Advances in Artificial Life: Proceedings of the 8th European Conference on Artificial Life*, volume 3630 of *Lecture Notes in Artificial Intelligence*, pages 744–753. Springer.

Klyubin, A. S., Polani, D., and Nehaniv, C. L. (2005b). Empowerment: A universal agent-centric measure of control. In *Proceedings of the 2005 IEEE Congress on Evolutionary Computation*, volume 1, pages 128–135. IEEE Press.

Prokopenko, M., Gerasimov, V., and Tanev, I. (2006). Evolving spatiotemporal coordination in a modular robotic system. In Nolfi, S., Baldassarre, G., Calabretta, R., Hallam, J., Marocco, D., Meyer, J.-A., and Parisi, D., editors, *From Animals to Animats 9: 9th International Conference on the Simulation of Adaptive Behavior (SAB 2006), Rome, Italy, September 25-29 2006*, volume 4095 of *Lecture Notes in Computer Science*, pages 558–569. Springer.

Sabidussi, G. (1966). The centrality index of a graph. *Psychometrika*, 31(4):581–603.

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

Vergassola, M., Villermaux, E., and Shraiman, B. I. (2007). 'infotaxis' as a strategy for searching without gradients. *Nature*, 445(7126):406–409.

Wasserman, S. and Faust, K. (1994). *Social Network Analysis*. Cambridge University Press.

Şimşek, Ö . and Barto, A. (2007). Betweenness centrality as a basis for forming skills. Technical report, University of Massachusetts, Department of Computer Science.