Evolving Computational Dynamical Systems to Recognise Abnormal Human Motor Function

Michael A. Lones[1], Stephen L. Smith[1], Andy M. Tyrrell[1], Jane E. Alty[2], and D. R. Stuart Jamieson[2]

[1] Department of Electronics, University of York, Heslington, York YO10 5DD, UK
{mal503,sls,amt}@ohm.york.ac.uk
[2] Leeds General Infirmary, Leeds, LS1 3EX, UK

Abstract. Artificial biochemical networks (ABNs) are a class of computational automata whose architectures are motivated by the organisation of genetic and metabolic networks. In this work, we investigate whether evolved ABNs can carry out classification when stimulated with time series data collected from human subjects with and without Parkinson's disease. The evolved ABNs have accuracies in the region of 80-90%, significantly higher than the diagnostic accuracies typically found in initial clinical diagnosis. We also show that relatively simple ABNs, comprising only a small number of discrete maps, are able to recognise the abnormal patterns of motor function associated with Parkinson's disease.

1 Introduction

We recently developed a series of computational dynamical systems motivated by the structure and function of biochemical networks [7, 8]. When evolved using an evolutionary algorithm, these *artificial biochemical networks* were shown to be competent at solving a diverse range of difficult control problems. In this work, we investigate whether artificial biochemical networks can be used to solve a difficult classification task, by distinguishing between movement time series data collected from Parkinson's disease patients and age-matched controls in a recent clinical study. In particular, we look at the classification accuracy of artificial biochemical networks which are composed of only a few non-linear discrete maps. Discrete maps, such as the logistic map and Chirikov's standard map, model complex real-world processes using simple iterative equations. Their dynamics make them computationally interesting in their own right, and when coupled together they have been shown to carry out difficult computational tasks [1].

2 Materials and Methods

2.1 Data Sets

Parkinson's Disease (PD) is a chronic neuro-degenerative disorder caused by the loss of dopamine-generating cells in the brain. The symptoms of PD are

highly variable, but all patients develop some form of movement disorder—such as slowing of movement (*bradykinesia*), tremor, rigidity, and impaired balance. Because of its symptomatic diversity, and symptom overlap with other diseases, PD is sometimes difficult to diagnose, with clinical misdiagnosis rates in the region of 25% [2, 5].

In a recent clinical study, we collected movement data from 49 PD patients and 41 age-matched controls as they performed a finger tapping task, a standard clinical means of measuring bradykinesia. The subject was asked to tap their thumb and index finger repeatedly for a duration of 30 seconds, using each hand in turn. Movement data was collected using a Polhemus Patriot electromagnetic motion tracking device, whose probes were attached to the subject's thumb and index finger whilst carrying out the task. The movement data was then transformed into time series of displacements between thumb and index finger, and divided into training and test sets in the ratio 2:1. The training data was used for fitness evaluation, and the test set was used to measure classifier generality.

2.2 Classifier Architectures

Artificial biochemical networks (ABNs) are a class of computational automata whose form and function are loosely modelled upon the biochemical networks found within biological organisms. In [7] and [8] we developed various ABN models, and showed how they display rich computational behaviours when coupled to a spectrum of dynamical systems. In this work, we are interested in whether ABNs can perform classification when stimulated with patient movement data. This approach is based on the hypothesis that ABNs can be evolved which will react to the dynamics found within a movement time series, causing them to alter their internal state in an observable manner. The approach is comparable to other uses of computational dynamical systems to perform time series classification, for example recurrent neural networks [3].

In this work, we use two types of ABN: artificial metabolic networks (AMNs) and artificial genetic networks (AGNs). The former are loosely modelled upon metabolic networks, and involve a set of functional elements (termed *enzymes*) manipulating an indexed set of real numbers (*chemical concentrations*) over a period of iterations. In order to mimic the effect of mass conservation in biology, the sum of chemical concentrations are normalised after each iteration. AGNs are a model of genetic regulation, and comprise an indexed array of *genes*, each of which has a real-valued state (*expression level*), a function (*regulatory function*) and a set of inputs (indices of other genes). The AGNs are synchronously updated over a period of iterations, with each gene's expression level at each iteration determined by applying its regulatory function to the expression levels of its input genes. Enzyme and regulatory functions are selected from a set of non-linear discrete maps: the logistic map, Chirikov's standard map, the baker's map, and Arnold's cat map [8]. Between them, these maps display a wide range of ordered and chaotic dynamical behaviours.

Inputs are delivered to the ABNs by setting a chemical concentration or the expression level of a designated input gene. Outputs are read from the final state

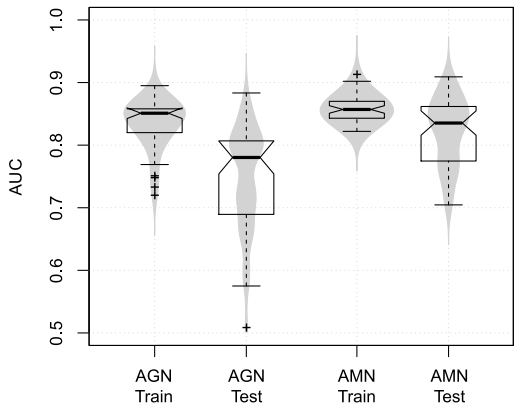

Fig. 1: Diagnostic power of evolved ABNs on both the training and test sets. Notched box plots show summary statistics over 50 runs. Overlapping notches indicate when median values (thick horizontal bars) are not significantly different at the 95% confidence level. Kernel density estimates of underlying distributions are also shown (in grey).

of a designated chemical concentration or gene expression level. A time series is delivered to a network one value at a time, each followed by t_b iterations of the network. Once the whole time series has been delivered, the network is executed for another t_a iterations in order to allow the dynamics to settle. At this point a single output value is read, which is then interpreted as the network's classification of the sequence. The settling parameters, t_b and t_a are both evolved with the network.

2.3 Evolutionary Algorithm

We used a standard generational evolutionary algorithm to evolve ABN-based classifiers. This used tournament selection of size 4, a single elite, a point mutation rate of 6% and uniform crossover with a crossover probability of 15%. Initial solution sizes were made intentionally short, between 2 and 10 genes/enzymes, to encourage parsimony. Each evolutionary run had a population size of 200 and a generation limit of 100. The fitness function was the area under the ROC curve (AUC), which is equivalent to the probability that a network will generate a higher output value for a PD time series than one from a control subject. Its relationship to probability means that AUC is easy to interpret, making it a popular metric in medicine [4].

3 Results

Fig. 1 shows the distribution of training and test scores for the best classifiers from each of 50 evolutionary runs, showing that AUC scores approach 0.9 for the

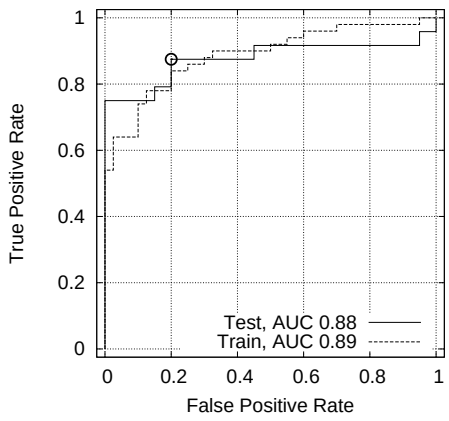

Fig. 2: ROC curves showing the diagnostic accuracy of an evolved classifier across all thresholds. The best trade-off is marked with a circle, showing a sensitivity of 87.5% and a specificity of 80%.

best parameter sets. This corresponds to classification accuracies in the range of 80-90% (see Fig. 2 for an example ROC curve). Whilst lower than the 92-94% accuracy of diagnosis performed by experts in movement disorders, it is considerably higher than the diagnostic accuracies found in non-expert secondary care (75%) and community care (47%), where most patients are first diagnosed [9]. This level of misdiagnosis led to the UK's National Institute of Clinical Excellence (NICE) to recommend that suspected PD patients should only be diagnosed by specialists. However, with this level of diagnostic accuracy, it is feasible that these kinds of classifiers could be used to assist primary care professionals such as general practitioners—especially given the relatively low cost of the equipment and the non-invasive nature of the diagnostic process.

Unlike our earlier diagnostic work [6, 10], which used window-based GP classifiers, ABNs have access to both local (e.g. local patterns of acceleration) and global (e.g. spectral characteristics) features of the data, so in principle are able to base their classification on diverse factors. This seems particularly important when processing movement data from Parkinson's patients, where symptom diversity means that individual indicators have poor diagnostic accuracy. For instance, the presence of tremor (measured using spectral analysis of at-rest data collected during the same clinical trial) has a diagnostic accuracy of 63%.

Our results are also interesting from an information processing perspective. Whilst the best classifiers generally make use of several discrete maps, good classification accuracy can be achieved with networks containing only one or two functional elements. Fig. 3 gives an example of this, showing the behaviour of an AMN containing only a single discrete map—Chirikov's standard map operating within a majority chaotic phase. It is surprising that a single chaotic map, in concert with the mass conservation rule of the AMN, can achieve a relatively

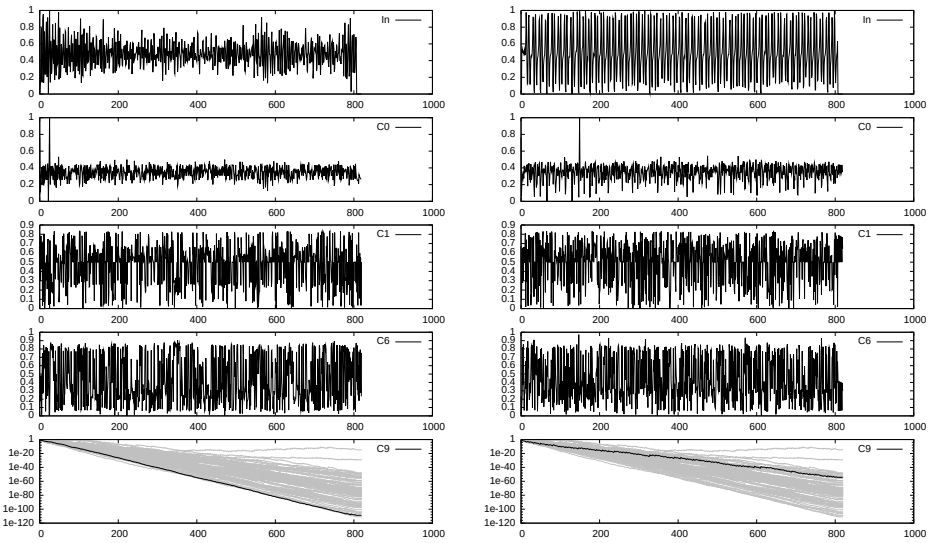

(a) Parkinson's disease patient (b) Control subject

Fig. 3: Example of an evolved AMN processing movement data (In) from subjects with and without Parkinson's. The lower four plots in each case show how chemical concentrations change as the input sequence is processed. C9 is the designated output chemical, whose final concentration is interpreted as the classifier's output. For comparison, grey lines show the C9 time series for all members of the test set.

high classification accuracy (AUC=0.84). Furthermore, this capacity does not appear to be linked to the choice of discrete map, since solutions with similar classification accuracies were evolved which contained each of the discrete maps available to the evolutionary algorithm.

Given the gap between ABNs and the biological structures which they are motivated by, it is hard to say whether these results give any insight into the nature of biochemical information processing. However, it does show that relatively simple (from an implementation perspective) non-linear processes can process signals produced by a relatively complex biological process. These non-linear processes, in turn, occur in many naturally occurring systems, so it doesn't seem unreasonable that they could occur in the biochemical networks present within cells and tissues.

4 Conclusions

In this paper, we have shown that artificial biochemical networks can be used to recognise abnormal motor function associated with Parkinson's disease. The evolved classifiers perform an objective diagnosis based upon data collected from simple movement tasks, and have accuracies comparable to trained clinicians.

Analysis of the classifiers suggests that diagnosis can be performed by relatively simple evolved networks, and that chaotic dynamics may play an interesting role. In future work, we hope to investigate whether this approach can also be applied to other forms of neurological disorder, such as Alzheimer's and Huntington's disease.

Acknowledgements

This work was funded, in part, by the EPSRC grant "Artificial Biochemical Networks: Computational Models and Architectures" (ref: EP/F060041/1).

References

1. Andersson, C., Nordahl, M.: Evolving coupled map lattices for computation. In: Banzhaf, W., et al. (eds.) Genetic Programming, Proc. 1st European Workshop on Genetic Programming, Lecture Notes in Computer Science, vol. 1391, pp. 151–162. Springer (1998)
2. Bajaj, N.P.S., Gontu, V., Birchall, J., Patterson, J., Grosset, D.G., Lees, A.J.: Accuracy of clinical diagnosis in tremulous parkinsonian patients: a blinded video study. J Neurol Neurosurg Psychiatry 81(11), 1223–1228 (Nov 2010)
3. Hüsken, M., Stagge, P.: Recurrent neural networks for time series classification. Neurocomputing 50, 223–235 (2003)
4. Kraemer, H.C., Morgan, G.A., Leech, N.L., Gliner, J.A., Vaske, J.J., Harmon, R.J.: Measures of clinical significance. J Am Acad Child Adolesc Psychiatry 42(12), 1524–1529 (Dec 2003)
5. Levine, C.B., Fahrbach, K.R., Siderowf, A.D., Estok, R.P., Ludensky, V.M., Ross, S.D.: Diagnosis and treatment of Parkinson's disease: a systematic review of the literature. Evid Rep Technol Assess (Summ) (57), 1–4 (May 2003)
6. Lones, M.A., Smith, S.L.: Objective assessment of visuo-spatial ability using implicit context representation cartesian genetic programming. In: Smith, S.L., Cagnoni, S. (eds.) Genetic and Evolutionary Computation: Medical Applications. John Wiley & Sons, Chichester, UK (2010)
7. Lones, M.A., Tyrrell, A.M., Stepney, S., Caves, L.S.D.: Controlling complex dynamics with artificial biochemical networks. In: Esparcia-Alczar, A.I., et al. (eds.) Proc. 2010 European Conference on Genetic Programming (EuroGP 2010). Lecture Notes in Computer Science, vol. 6021, pp. 159–170. Springer Berlin / Heidelberg (2010)
8. Lones, M.A., Tyrrell, A.M., Stepney, S., Caves, L.S.D.: Controlling legged robots with coupled artificial biochemical networks. In: Lenaerts, T., et al. (eds.) Advances in Artificial Life, ECAL 2011: Proc. 11th European Conference on the Synthesis and Simulation of Living Systems. pp. 465–472. MIT Press (August 2011)
9. National Institute for Health and Clinical Excellence: Parkinson's disease: diagnosis and management in primary and secondary care. Royal College of Physicians (2006), http://www.nice.org.uk/CG035
10. Smith, S.L., Gaughan, P., Halliday, D.M., Ju, Q., Aly, N.M., Playfer, J.R.: Diagnosis of Parkinson's disease using evolutionary algorithms. Genetic Programming and Evolvable Machines 8(4), 433–447 (2007)