

Deep Learning Applied to Attractor Images Derived from ECG Signals for Detection of Genetic Mutation

Philip J Aston^{1,2}, Jane V Lyle¹, Esther Bonet-Luz¹, Christopher LH Huang³, Yanmin Zhang⁴, Kamalan Jeevaratnam¹, Manasi Nandi⁵

¹ University of Surrey, Guildford, UK

² National Physical Laboratory, Teddington, UK

³ University of Cambridge, Cambridge, UK

⁴ Affiliate Children's Hospital of Xi'an Jiaotong University, Xi'an, China

⁵ King's College London, London, UK

Abstract

The aim of this work is to distinguish between wild-type mice and Scn5a^{+/-} mutant mice using short ECG signals. This mutation results in impaired cardiac sodium channel function and is associated with increased ventricular arrhythmogenic risk which can result in sudden cardiac death. Lead I and Lead II ECG signals from wild-type and Scn5a^{+/-} mice are used and the mice are also grouped as female/male and young/old.

We use our novel Symmetric Projection Attractor Reconstruction (SPAR) method to generate an attractor from the ECG signal using all of the available waveform data. We have previously manually extracted a variety of quantitative measures from the attractor and used machine learning to classify each animal as either wild-type or mutant. In this work, we take the attractor images and use these as input to a deep learning algorithm in order to perform the same classification. As there is only data available from 42 mice, we use a transfer learning approach in which a network that has been pretrained on millions of images is used as a starting point and the last few layers are changed in order to fine tune the network for the attractor images.

The results for the transfer learning approach are not as good as for the manual features, which is not too surprising as the networks have not been trained on attractor images. However, this approach shows the potential for using deep learning for classification of attractor images.

1. Introduction

Cardiac arrhythmias, arising from cardiac ion channel defects, can be complex to diagnose and often require detailed 12 lead ECG analysis focussing on intervals and amplitudes. Our approach is to use the entire ECG waveform shape, rather than specific intervals or amplitudes, as this

may reveal differences that are less apparent in conventional ECG analysis on the more commonly used Lead I and Lead II signals.

Experimental studies often use mouse models as they are amenable to genetic modification. We consider short ECG signals from wild-type and Scn5a^{+/-} mutant mice. This mutation is a model for Brugada Syndrome, which is associated with sudden cardiac death [1].

We have previously applied the novel Symmetric Projection Attractor Reconstruction (SPAR) method [2, 3] to this mouse data [4], which involved generating an attractor from the ECG signals from which features were manually extracted and machine learning was then used to classify the signals as being from either a wild-type or mutant mouse. In this work, we use deep learning applied to the attractor images directly in order to perform the same classification task. In particular, we modify pre-trained deep neural networks to give our binary classification, which are then retrained on the attractor images, a process known as transfer learning. A similar approach has been applied to other biological image classification problems, including detecting metastases in whole slide images of lymph node sections [5], brain abnormality classification using magnetic resonance images [6] and alcoholism screening of brain magnetic resonance images [7].

2. The Data

Short ECG signals were obtained from 42 anaesthetised mice using the procedure described in [8]. Lead I and II recordings were obtained from 36 mice with an additional 3 Lead I and 3 Lead II recordings from a further 6 animals, giving 78 signals in total with duration ranging from 21 seconds to 10 minutes. The 78 signals can be classified by type (42 WT, 36 Scn5a^{+/-}), gender (42 male, 36 female) and age (39 young, 39 old).




Figure 1. The SPAR method applied to an ECG signal (from [9]). From left to right: (i) Lead II ECG signal; (ii) Three-dimensional attractor obtained using Takens' delay coordinates; (iii) Two-dimensional attractor obtained by projecting the three-dimensional attractor onto a plane perpendicular to the vector $(1,1,1)$; (iv) The two-dimensional attractor density.

3. The Symmetric Projection Attractor Reconstruction Method

Classical analysis of ECG signals involves identifying particular points on the signal (P, Q, R, S, T) from which a variety of intervals, such as QT, ST, PR, etc., can be derived. Some amplitudes can also be found but are less commonly used. However, this approach derives only a small number of features from the large quantity of data points captured by an ECG device. Our SPAR method uses all of the available waveform data and repackages it in the form of a two-dimensional attractor, using Takens' delay coordinates together with some additional steps, and so the waveform shape is encoded in the shape of the attractor (see Fig. 1) [2,3]. Clearly, any changes in the shape of the waveform will result in corresponding changes in the morphology of the attractor.

We have previously applied the SPAR method to the analysis of this mouse data [4]. In this case, we manually extracted 74 features from each attractor to which we applied machine learning (k nearest neighbours ($k = 3$) with forward feature selection) to provide the classification. The manual features describe the geometry, density and symmetry of the attractor. We now consider an alternative approach in which we use deep neural networks for automatic feature extraction using convolutional neural networks, thus removing the requirement for the manual feature extraction. The results obtained using this deep learning approach are compared with the results obtained using manual feature extraction in Section 5.

4. Transfer Learning

There are 16 pre-trained deep neural networks available in the Deep Learning Toolbox of Matlab 2019a that have been trained on the ImageNet database [10] which contains over one million images classified into one thousand object categories. We considered only four of these networks, namely AlexNet, GoogLeNet, ResNet-18 and SqueezeNet. Transfer learning involves replacing the final classification layers of a pre-trained network and then

retraining it on a new set of images. In our case, we generate attractor images which are labelled as either wild-type (WT) or $Scn5a^{+/-}$ mutation (SCN) and these images are used to retrain the networks. All the weights and biases were frozen, except for the new classification layer and we used 20 epochs to retrain the network.

Following the same process as in [4], we take thirteen consecutive windows each of length 10 seconds from all of the ECG signals and generate an attractor for every window. The classification is performed for each window of data/attractor individually and then a majority vote is used to classify the animal. The advantage of using multiple windows is that more images are generated (1014 in total), but clearly the disadvantage is that, particularly for very stable signals, there is a lot of repetition in the images. However, we will use this approach as it allows for a simple comparison with our previous results.

The results that we report are from cross-validation. In this case, we divide the data into 18 groups, ensuring that all 26 images derived from Lead I and Lead II signals from a particular animal are contained in the same group, so that none of the data from animals in the validation set is contained in the training data. Each group contains an equal number of WT and SCN images in most cases, although this not possible in all cases due to the unequal size of these two groups. The cross-validation error is minimised using 100 iterations of Bayesian optimisation over the four parameters of learning rate, minibatch size, L2 regularisation and momentum.

We then repeat this procedure for classification of the signals according to age (young/old) and gender (male/female).

5. Results

We compare the results obtained for the three binary classifications described above using the transfer learning approach with the results that we previously obtained using manual features and feature selection [4]. A summary of these results, by individual record and by animal, is shown in the top row of Table 1. If the number of correctly classi-




Figure 2. Histograms showing the distribution of votes for the classification of Type for the four different pre-trained networks. Animals are correctly classified if the number of correct votes is 7 or above.

fied images from the 13 used for each animal is 7 or above then, by majority vote, the animal is correctly classified.

We note that none of the transfer learning results gave as high an accuracy as was achieved with the manual features, except for the classification of Gender by record using GoogLeNet. Type classification had the highest accuracy for manual features (85.9%) and for all deep learning models, with GoogLeNet again achieving the best deep learning result (79.8%). However, in Age classification, AlexNet provided the best deep learning result (70.8% against 79.2% for manual features), and a much larger deficit in accuracy was observed in the other models.

In terms of performance on the ImageNet dataset, of the four networks we are considering, the ResNet-18 network achieved the highest accuracy [11]. However, in our case, it was the worst performing for classification of Type and was never the best performing network. It is also interesting to note that the best performing network was not consistent across the different classifications. Overall, the performance of SqueezeNet was poor.

Histograms for each of the four networks showing the number of correct votes are shown in Fig. 2 for the classification of Type. The number of animals for which all 13 images were correctly classified is 51 for AlexNet, 45 for GoogLeNet, 28 for ResNet-18 and 42 for SqueezeNet. Thus, AlexNet scored 13/13 for almost two thirds of the animals. However, it also had the highest number of animals (11) where all 13 images were misclassified.

There were 20 animals that had all 13 images correctly classified by all four networks which consisted of 15 wild type and 5 mutant. Two attractors from this group are shown in Fig. 3. There was also 1 animal (SCN) that had all 13 images misclassified by all the networks. The 6 lowest total correct votes across all networks were all for SCN animals while the lowest total correct vote for a WT animal was 12/52. An attractor from each of these animals is shown in Fig. 4. The very thin arms of the WT attractor in Fig. 3 is very common for WT animals, and so it is easy to see why the SCN attractor in Fig. 4 was misclassified. Similarly, the variability shown in the WT attractor in Fig. 4 is very characteristic of SCN attractors.

For the 20 animals that scored 13/13 for all networks,

Table 1. Deep learning results for binary classifications compared with results using manual features. Top row: accuracy of individual records; bottom row: accuracy for animals using majority vote. The best/worst result for the neural networks in each column is coloured green/red.

	Type (WT/SCN)	Age (Young/Old)	Gender (M/F)
Manual Features	85.9% 68/78	79.2% 66/78	76.7% 65/78
AlexNet	77.1% 61/78	70.8% 57/78	68.2% 55/78
GoogLeNet	79.8% 66/78	64.8% 53/78	77.1% 62/78
ResNet-18	74.2% 62/78	61.5% 50/78	68.4% 54/78
SqueezeNet	76.1% 62/78	56.4% 45/78	58.6% 50/78

the classification with manual features also scored 13/13 in 19 cases with 1 scoring 12/13 and so the results here are consistent. However, it is interesting to note that for the animal in each class with the lowest total correct votes across all the networks, the classification with the manual features gave scores of 13/13 for the SCN animal and 6/13 for the WT animal and so clearly the manual features are quantifying aspects of the attractor that are missed by the feature extraction of the deep neural networks.

We conclude on a light-hearted note, by classifying the four attractors shown in Figs 3 and 4 into one of the 1,000 ImageNet object categories. The attractors were all classified using the four pre-trained networks as either “barn spider”, “walking stick”, “spotlight” or “harvestman” (also known as “daddy longlegs”). The associated probabilities are generally quite low (0.1–0.4), although ResNet-18 classified the WT attractor in Fig. 3 as a “spotlight” with probability of 0.77!

6. Conclusions

Clearly the pre-trained networks used for the transfer learning have not been trained on images that resemble our attractors, whereas the manual features have been tailored to the specific features of the attractors, and so it is not too

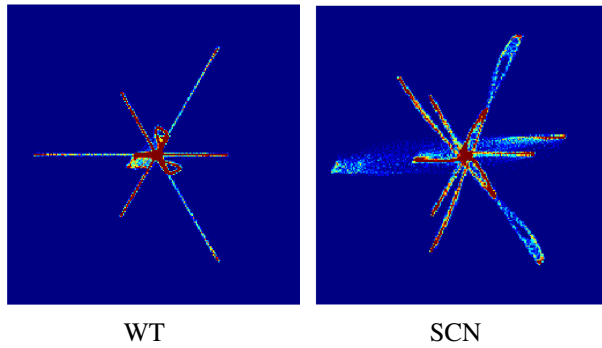


Figure 3. Attractors from a wild type (left) and a mutant (right) animal for which all four networks classified all 13 windows correctly.

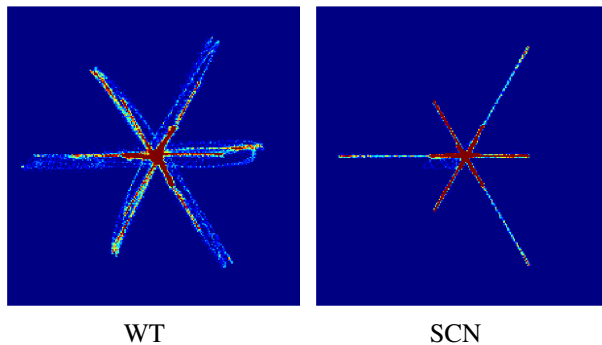


Figure 4. Attractors from a wild type (left) and a mutant (right) animal for which all four networks classified all 13 windows incorrectly.

surprising that they have not been able to match the prediction accuracy obtained with the manual features. The other issue is that our dataset is small and has repeated windows containing similar information, at least for stable signals. These two issues can be addressed by using a larger dataset and developing and training deep neural networks specifically for classification of attractor images and it is anticipated that this would give improved results.

We also note that the feature selection that we used with the manual features resulted in an improvement in accuracy. Performing classification using all 74 of the manual features without feature selection using the k nearest neighbours classifier gives accuracies of 71.8% for Type, 53.0% for Age and 56.9% for Gender. Clearly the transfer learning approach exceeds all of these. There is however no similar feature reduction process with deep learning.

This study has shown the potential for using deep neural networks for the classification of attractor images, but clearly further work is required.

Acknowledgments

The authors would like to thank the University of Surrey for access to its High Performance Computing facility. This project (18HLT07) has received funding from the EMPiR programme co-financed by the Participating States

and from the European Union’s Horizon 2020 research and innovation programme.

References

- [1] J. Brugada, R. Brugada, C. Antzelevitch, J. Towbin, K. Nademanee and P. Brugada, “Long-term follow-up of individuals with the electrocardiographic pattern of right bundle-branch block and ST-segment elevation in precordial leads V1 to V3,” *Circulation* vol. 105, pp. 73–78, 2002.
- [2] P.J. Aston, M.I. Christie, Y.H. Huang and M. Nandi, “Beyond HRV: Attractor reconstruction using the entire cardiovascular waveform data for novel feature extraction,” *Phys. Meas.* vol. 39, 024001, 2018.
- [3] M. Nandi, J. Venton and P.J. Aston, “A novel method to quantify arterial pulse waveform morphology: Attractor reconstruction for physiologists and clinicians,” *Phys. Meas.* vol. 39, 104008, 2018.
- [4] E. Bonet-Luz, J. V. Lyle, C. L. H. Huang, Y. Zhang, M. Nandi, K. Jeevaratnam and P.J. Aston, “Risk prediction of cardiac arrhythmogenic syndromes using Symmetric Projection Attractor Reconstruction analysis,” submitted.
- [5] T. Rai, A. Morisi, B. Bacci, N. J. Bacon, S.A. Thomas, R.M. La Ragione, M. Bober and K. Wells, “Can ImageNet feature maps be applied to small histopathological datasets for the classification of breast cancer metastatic tissue in whole slide images?” *Medical Imaging 2019: Digital Pathology*, eds J. E. Tomaszewski, A. D. Ward, Proc. of SPIE, vol. 10956, 109560V, 2019.
- [6] M. Talo, U. Baran Baloglu, Ö. Yildirim and U. Rajendra Acharya, “Application of deep transfer learning for automated brain abnormality classification using MR images,” *Cog. Sys. Res.* vol. 54, pp. 176–188, 2019.
- [7] S.-H. Wang, S. Xie, X. Chen, D. S. Guttery, C. Tang, J. Sun and Y.-D. Zhang, “Alcoholism identification based on an AlexNet transfer learning model,” *Front. Psychiatry* vol. 10, 205, 2019.
- [8] K. Jeevaratnam, Y. Zhang, L. Guzadhur, R. M. Duehmke, M. Lei, A. A. Grace and C. L. -H. Huang, “Differences in sino-atrial and atrio-ventricular function with age and sex attributable to the $Scn5a^{+/-}$ mutation in a murine cardiac model,” *Acta Physiol.* vol. 200, pp. 23–33, 2010.
- [9] J. V. Lyle, P. H. Charlton, E. Bonet-Luz, G. Chaffey, M. Christie, M. Nandi and P.J. Aston, “Beyond HRV: Analysis of ECG signals using attractor reconstruction,” *Computing in Cardiology* vol. 44, pp. 091–096, 2017.
- [10] ImageNet database: <http://www.image-net.org/>
- [11] Pre-trained deep neural networks. <https://uk.mathworks.com/help/deeplearning/ug/pretrained-convolutional-neural-networks.html>

Address for correspondence:

Prof Philip J Aston

Department of Mathematics, University of Surrey, Guildford, Surrey GU2 7XH, UK and

Data Science Group, National Physical Laboratory, Hampton Road, Teddington TW11 0LW, UK

P.Aston@surrey.ac.uk; Philip.Aston@npl.co.uk