

Transparency Of Execution Using Epigenetic Networks

Alexander Turner¹ and Nina Dethlefs¹

¹ Department of Computer Science, University of Hull, UK
alexander.turner@hull.ac.uk

Abstract

This paper describes how the recurrent connectionist architecture epiNet, which is capable of dynamically modifying its topology, is able to provide a form of transparent execution. EpiNet, which is inspired by eukaryotic gene regulation in nature, is able to break its own architecture down into sets of smaller interacting networks. This allows for autonomous complex task decomposition, and by analysing these smaller interacting networks, it is possible to provide a real world understanding of *why* specific decisions have been made. We expect this work to be useful in fields where the risk of improper decision making is high, such as medical simulations, diagnostics and financial modelling. To test this hypothesis we apply epiNet to two data sets within UCI's machine learning repository, each of which requires a specific set of behaviours to solve. We then perform analysis on the overall functionality of epiNet in order to deduce the underlying rules behind its functionality and in turn provide transparency of execution.

Introduction

Since their inception, machine learning techniques have been developed with a key focus on performing a task objectively better than other machine learning counterparts. Frequently machine learning techniques have even been able to outperform their human counterparts, a trend which will inevitably continue with the ever increasing computational resources available (He et al., 2015). Machine learning techniques, especially connectionist architectures such as neural networks (Lones et al., 2013), frequently function as 'black boxes' where it can be difficult to understand the rationale for their decision making process, which can be problematic in fields where the objective performance *must* be aligned with the current knowledge of the task. Take for example the field of medicine using convolutional neural networks for classification of various radiographic samples. Convolutional neural networks are the state of the art in object recognition (Szegedy et al., 2015) however they can be easily fooled and will confidently classify images which are unrecognisable when compared to the target image (Nguyen et al., 2015). In the medical field it is highly important to

know *why* a certain feature is selected to represent a diagnosis. Otherwise, it is possible that the classification is being achieved due to an artefact of the image, such as how well it is focused, rather than the relevant pictographic criteria to make a diagnosis (Nguyen et al., 2015). Such issues are also echoed in the use of artificial intelligence for control systems, specifically self-driving cars and applications in financial decision making.

The process of understanding *why* machine learning techniques make specific decisions is difficult due to the complexity of their architecture. EpiNet is a machine learning technique inspired by the functionality of chromatin remodelling in nature which allows for the selective expression and repression of genes according to environmental stimulus. In nature every gene within an organism has the potential to be expressed, but only certain genes are expressed at certain time points. By looking at what genes are being expressed, it is often possible to understand why an organism is behaving as it is (Smith et al., 2014; Manshian et al., 2015), without having to understand the entire functionality of the organism. EpiNet is similar in this sense, as by analysing only the expressed genes, and how they interact with a task, it is possible to deduce rules for that section of expressed genes. This effectively allows the reconstruction of a whole networks' behaviour via the analysis of the significantly smaller sub-sections of expressed genes.

This paper builds upon this premise by applying epiNet to two different tasks, to understand the benefits and difficulties when trying to provide transparency of execution. That is, the ability to query its behaviour and have a definitive reason as to why that behaviour exists via the statistical analysis of the dynamics of the network.

Gene Regulation and Epigenetic Processes

DNA is ubiquitous in the natural world, and is used by nature to store the particular information about an organism in which it resides. A gene can be considered a subsection of DNA which is commonly used to describe the structure of a protein, which in turn can be considered a small molecular machine used in many processes, most notably to build

tissue, break down and metabolise energy sources and in stimulus response. For a gene to become expressed, cellular machinery must bind to the DNA strand where that gene is located, and transcribe that particular section of DNA into a complimentary RNA strand. This RNA strand, which holds the information from the gene it was just transcribed from, is used as a template which forms the primary structure of an amino acid sequence. This sequence can then fold to become a functional protein. To move from DNA strand to a protein requires the accurate binding of around 20 interacting proteins, which in turn are products of other genes. Hence, within human cells which have around 20000 genes, the regulatory map, depicting the interactions between genes is complex Clark and Pazdernik (2013). There are two key properties of gene regulatory networks. The first is that they are dynamic, where connections between genes may only be present for a limited time. The second is that apart from a few exceptions such as red blood cells, all cells contain the full set of regulatory genes, and can modify which genes are active out of this set to allow the differentiation into various cell types. Changes in the subset of active genes can modify cells' fate. This is in combination with the other 'housekeeping' genes which are pervasive in almost all cell types. Both long term changes such as cellular differentiation and short term changes such as a stress response can both be achieved by selecting, out of all possible genes, which ones are to be active at any given time.

Epigenetics

Epigenetics can be described as stable modifications of gene expression without alteration of the genetic code. More colloquially epigenetic processes can modify the expression of genes by interrupting, most commonly, the process of transcription and translation. Two of the most ubiquitous epigenetic processes are chromatin remodelling and DNA methylation.

Chromatin Remodelling Chromatin is the higher order folding of nucleosomes, which is the combination of DNA and histone proteins (Figure 1). Chromatin serves two key purposes. Firstly it provides the ability for long stands of DNA to be effectively compressed and be able to fit in the cells relatively small nucleus. Secondly, by changing the structure of chromatin over time, it allows the cellular machinery to access different parts of the genome, and in turn express different genes. This is one of the principle methods of gene regulation in eukaryotic cells Gentry and Hennig (2014).

DNA Methylation DNA methylation is the addition of a methyl base to DNA which acts as a physical barrier preventing the binding of transcription complexes to the DNA. Whereas chromatin remodelling can be thought of as facilitating more short term changes in gene expression, DNA

methylation can be thought of as long term repression of gene expression. DNA methylation is pervasive in nature, with species having between 0.0002 and 14% of their DNA bases methylated Zemach et al. (2010); Capuano et al. (2014).

EpiNet

EpiNet is the most recent incarnation of an epigenetic network, which differs from its previous counterparts Turner et al. (2013a,b) because it allows a *dynamic* selection of genes to be executed at every time step. EpiNet comprises of a set of nodes called genes which are abstracted from their biological counterparts. Each artificial gene exists within a space called the reference space, where if genes overlap within this space they are connected. In terms of each genes' regulatory dynamics, they contain a parameterisable sigmoid function, where the parameters are stored by the genes themselves, hence each gene can have varying regulatory functions. The inputs to the gene are taken from the expressions and weights of the genes that are connected to it. The result of this function is then that gene's expression for the current time step. This collection of genes is known as a gene regulatory network (GRN) and is the backbone of epiNet. The GRN by itself is a valid machine learning technique capable of solving complex tasks and can be described as follows :

G is a set of genes $\{n_0 \dots g_{|G|} : g_i = \langle a_i, I_i, W_i \rangle\}$ where:

$a_i : R$ is the activation level of the node.

$I_i \subseteq G$ is the set of inputs used by the node.

W_i is a set of weights, where $0 \leq w_i \leq 1, |W_i| = |I_i|$.

L is a set of initial activation levels, where $|L_G| = |G|$.

$In \subset G$ is the set of nodes used as external inputs.

$Out \subset G$ is the set of nodes used as external outputs.

The GRN architecture is similar to an recurrent neural network, with the exception that the GRN allows the genes (nodes in the recurrent neural network) to contain a range of parameterisable regulatory functions. The connections between the environmental inputs and outputs from the GRN are static throughout execution. The epigenetic molecules within epiNet sit 'on top' of a GRN, akin to histone proteins sitting 'on top' of DNA in nature. The epigenetic molecules purpose is to dynamically and autonomously select genes from the GRN for execution according to environmental inputs. Each epigenetic molecule is connected to a subset of available genes (via its position within the reference space), and at each time step it takes the expressions from these genes and processes them within its own regulatory function. This value is then used to move the epigenetic molecule within the reference space, which updates which genes are selected for execution in the next time step. Additionally, it also modifies which genes are connected to the epigenetic molecule. This process is repeated at each time step. The

inputs and outputs are then mapped to and from the genes which have been selected via the epigenetic molecules. An image showing the execution of the epigenetic network can be seen in Figure 1. The epigenetic molecules can be formally described as follows :

E is a set of epigenetic molecules
 $\{e_0 \dots s_{|E|} : e_i = \langle a_i, I_i, W_i, C_i \rangle\}$:
 $a_i^e \in R$ is the position of the molecule.
 $I_i^e \subseteq N$ is the set of inputs to the molecule.
 W_i^e is a set of weights, where $0 \leq w_i \leq 1, |W_i| = |I_i|$.
 $C_i^e \subseteq N$ is the set of nodes controlled by the switch.

The algorithm describing the application of epiNet to a task can be seen in algorithm 1.

Algorithm 1 Evaluating epiNet on a task

```

1: initialize control task
2:  $a \leftarrow L$  ▷ initialize epiNet state
3: repeat
4:    $cout \leftarrow$  state variables from controlled system
5:    $ln \leftarrow$  SCALE( $cout$ ) ▷ scale inputs to [0, 1]
6:   for  $i \in \{0, \dots, |E|\}$  do ▷ update positions of
     epigenetic molecules
7:      $a_i^s \leftarrow$  SIGMOID( $I_i^s \cdot W_i^s$ ) ▷ modify epigenetic
     positions
8:   end for
9:   for  $i \in \{0, \dots, |G|\}$  do ▷ update list of expressed
     genes
10:     $a_i \leftarrow$  SIGMOID( $I_i \cdot W_i$ ) ▷ execute expressed
    genes
11:  end for
12:   $cin \leftarrow$  SCALE(Out) ▷ scale outputs to range
13:  modify controlled system according to  $cin$ 
14: until control task finished or timed-out
15:  $fitness \leftarrow$  progress on control task objectives

```

Task Definitions and Application

The dynamical topological changes during the execution of epiNet mean that the classical optimisation techniques of artificial neural networks such as back propagation are not feasible for optimising epiNet. The way epiNet is optimised is to modify its static structure, that is the genes and epigenetic molecules and the data held within them, as the dynamic functionality of the network is a direct product of its static topology. Hence, to optimise epiNet, evolutionary algorithms or evolutionary strategies are the most well suited candidates. Typically this would be done with evolutionary algorithms, which are a population based optimisation technique. However, with the focus of this would being on understanding the functionality of epiNet to provide transparency, not on objective performance, a 1 + 1 mutation only

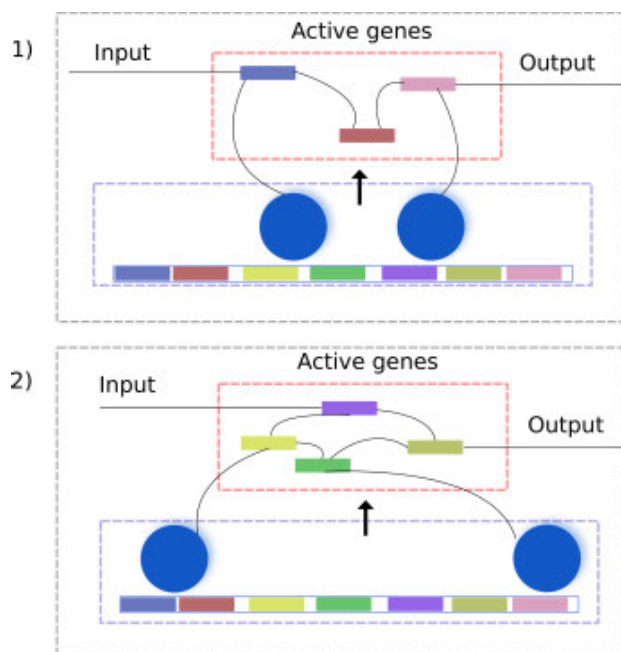


Figure 1: An abstraction from epiNet showing its execution over two consecutive time steps with only one input and one output being used. In the first time step, it can be seen that the genes which are not covered by the epigenetic molecules are selected for execution. The genes are then moved to a separate network where they self-organise and connect to each other and the inputs and outputs from the task. Once executed, the genes then feedback into the epigenetic molecules which use this information to update their position. In the second time step, the epigenetic molecules have moved and cover a different subsection of genes, allowing a different set of genes to be executed. This process is repeated until completion or cessation of the task.

evolutionary algorithm will be used as it will be possible to track the changes to the epiNet as it evolves, which will help provide insight of its functionality.

To assess epiNet’s ability to provide transparency during evolution it will be applied to two time series based classification tasks which are available from the University of California, Irvines machine learning repository (University of California, Irvine, School of Information and Computer Sciences, 2017). Each task will be conducted over 100000 evaluations of epiNet, where each data element of the network will be subjected to mutation at a probability of 5% at each iteration. An algorithm depicting the optimisation of epiNet can be seen in Algorithm 2.

Using Electroencephalogram Readings to Determine Eye State

Electroencephalogram (EEG) is a method of monitoring the electrical activity of the brain, which is typically achieved

Algorithm 2 Optimising epiNet

```

1: bestNetwork ← bestNetwork ∪ {new random
   epiNet}
2: for numberofevaluations do
3:   CLONE P AS BESTNETWORK
4:   MUTATE(P)
5:   EVALUATE(P) ▷ see Algorithm 1
6:   if P.fitness ≥ bestNetwork.fitness then
7:     bestNetwork = P
8:   end if
9: end for
  
```

non-invasively by attaching electrodes to the scalp. It is most commonly used to diagnose epilepsy (Parmeggiani et al., 2010), but is also finding significant traction as a brain-computer interface (Chambayil et al., 2010; Wairagkar et al., 2016).

The data for this task (University of California, Irvine, School of Information and Computer Sciences, 2017) was generated using the Emotiv EEG Neuroheadset (Emotive, 2017) to understand if analysis of the EEG signals can predict if the subject is blinking or not. 14 individual electrodes are connected to the subject and the EEG signals are recorded for 117.03 seconds at 128hz, resulting in 14980 data points. For this task, epiNet will be randomly initialised with 25 genes and 2 epigenetic molecules. The network will produce a single output which will be mapped to 0 if its value is less than 0.5 and 1 otherwise. For each time step, epiNet will try and deduce whether or not the subject is blinking or not. The data will be split into 70% for training and 30% for testing. The fitness for the network will be the normalised value of time steps which have been correctly classified. A total of 50 runs for the experiment will be conducted.

Ultrasound Data Of a Wall Following Robot

This data set contains the readings from 24 ultrasound sensors connected to a traversing robot. The robot has been programmed to perform a wall following task. The task took approximately 10 minutes with a 9hz sampling producing 5456 data points. The data set has 5 possible actions attached to each time step which are turn left, turn sharp left, turn right, turn sharp right and move forward. In this work, we have simplified this to 3 possible actions which are turn left, turn right and move forward, and used the compact data set attached to this data, which maps the 24 variables onto 4 variables, specifying the distance from the robot to the wall in each direction. The heuristic for this can be seen in Algorithm 3. When epiNet is applied to this task (Algorithm 1 and 2) it will produce 3 outputs, corresponding to each of the three possible actions. The output with the highest values will be taken as the decision at that time step. The data will be split into 70% for training and 30% for testing. The

Start	End	Length	Topology	Inputs	Output
1020	1100	70	1,3,4,5,7	1,4	0.3
1101	1102	1	1,3,4,5	1	0.1
1106	1304	198	1,2,3,4,5,7	1,4	0.9
1305	1400	95	1,3,4,5,7	1,4	0.3
1401	1402	1	1,3,4,5	1	0.1
1406	1604	198	1,2,3,4,5,7	1,4	0.9

Table 1: An example of a summary of the information autonomously collected and made available throughout the execution of epiNet. This information can be generated and queried live or post execution. The topologies column list which genes are expressed at any given time step, the inputs column shows which inputs from the tasks are currently connected to an expressed gene and the outputs shows the current output of the network.

fitness for the network will be the normalised value of time steps which have been correctly classified. A total of 50 runs for the experiment will be conducted.

Algorithm 3 Wall following heuristic

```

1: if Left distance > 0.9 then
2:   if Front distance ; 0.9 then
3:     Turn to the right
4:   else
5:     Turn to the left
6:   end if
7: else
8:   if Front distance ; 0.9 then
9:     Turn right
10:  else if Left distance ; 0.55 then
11:    Turn right
12:  else
13:    Move forward
14:  end if
15: end if
  
```

Autonomous Analysis of Network Function

The focus of this work is to provide an efficient way to analyse the dynamics of eipNet to provide a real world understanding of *why* it functions the way it does. To achieve this, we will be looking at specific data available to epiNet, an example of which can be seen in Table 1. The most relevant data in the table is the topological changes over time. With only a small proportion of genes being selected for execution at each time step by the epigenetic molecules, it is relatively simple to understand a specific topology, and also, how dynamical patterns of topological changes change according to the dynamics of the task. Additionally, algorithms specifically for analysing epiNet have been developed which can detect repeating patterns of topological change and effec-

tively condense this information so that it can be easily understood. In previous work, it has been noted how topological changes to the network frequently correlate with changes in the dynamics of the task. In addition to this, we can also see which inputs from the task are being used at any given time. For example, in the wall following task it might be possible to show that the majority of decision making by epiNet is based upon the information from only a small subsection of possible inputs. This type of information is useful for two key reasons. Firstly, it can be used to inform further scientific query. Secondly, it provides an understanding of how the network behaves by process of elimination. For example, when applied to the coupled inverted pendulums it was found that only two of the ten inputs from the task were needed to solve it optimally (Turner et al., 2017). This not only highlighted that only two inputs were generally used to make decisions within epiNet, but that the remaining 8 inputs served no purpose throughout optimisation. A further piece of information which is available through epiNet is the trace of each epigenetic molecule detailing its position over time. This serves to highlight the stability of specific dynamical regimes or topologies, and how likely changes to the position of the epigenetic molecules translate to changes in topology and network output.

Within this work, we are able to provide a look up table which stores all the information regarding epiNets execution at every time step. This means, that during a period of interest, say a changing of behaviours in the network, all the information of epiNets execution during that time can be efficiently extracted. In addition a summary can be provided of the overall execution of the network, showing the average size of each topology, how many and what inputs are most frequently used and which patterns of execution have been found.

Results

The focus of this work is on providing transparency during the execution of epiNet, however it is important to ascertain the computational potential of epiNet, as without this it is of limited use from a computational perspective. To do this we have run an identical set of experiments on the GRN which underpins EpiNet. Because EpiNet is built upon the GRN, which is a valid computational model in it's own right, the difference in performance can be attributed to the addition of the epigenetic molecules within EpiNet, as the underlying GRN will be identical to the one tested here. Figures 2 and 3 show that epiNet performs significantly better than the GRN architecture which it was based upon. The difference between performance was most pronounced in the wall following robot task, where epiNet was able to score maximum fitness when applied to the testing data. However, in this task the spread of results between the 50 runs of epiNet was much greater than the GRN architecture, which suggests that epiNet may be prone to getting stuck with local




Figure 2: The overall performances of the GRN architecture and epiNet on the EEG classification task. EpiNet significantly outperforms the GRN on both training and test data with significance values of 1.3123×10^{-15} and 6.8560×10^{-5} respectively, using the Wilcoxon rank-sum test.




Figure 3: The overall performances of the GRN architecture and epiNet on the robot wall following task. EpiNet significantly outperforms the GRN on both training and test data with significance values of 7.1311×10^{-13} and 6.0915×10^{-12} using the Wilcoxon rank-sum test.

optima. In terms of the EEG classification task, it was substantially more difficult for the networks to learn, with many instances of both networks performing worse than random on the training data. This highlights that a type of cross fold validation might be beneficial, however, it would have to be adapted due to how epiNet executes and is optimised.

Deducing Rules From The ‘Using Electroencephalogram Readings to Determine Eye State’ Task

To understand the process of providing transparency of execution and its usefulness, we will focus on a single network, the optimum performing epiNet according to the test data, achieving an 82% classification rate. This network in its original form contains 24 genes and 7 epigenetic molecules, significantly more than the 2 it contained at the beginning of the optimisation process. Within the scope of this paper, it will not be possible to explain every decision the network

Algorithm 4 Transparency Algorithm

- 1: Isolate time steps of interest
- 2: **for** $N \neq 1000$ **do**
- 3: Remove random gene / epi
- 4: Re-evaluate network
- 5: **if** !Full functionality is preserved **then**
- 6: Replace gene / epi
- 7: **end if**
- 8: **end for**
- 9: Analyse topology changes associated with changes in decision
- 10: Correlate which epigenetic molecules are responsible for the change
- 11: Calculate which inputs / genes key epigenetic molecules are connected to

made, hence we will focus on the one outlined in Figure 4 as the decision of interest. This decision made by epiNet is to predict that from the data, the subject had their eyes closed, and opened them at the 1200th time step. This decision is incorrect, as illustrated by Figure 4, as the subjects eyes remain closed until the 1590th time step. We will answer *why* this incorrect decision was made.

The network we are focusing on, with 24 genes and 7 epigenetic molecules is far too complex to deduce its functionality in its current state. The network was optimised to function over thousands of time steps, changing its topology depending on the network dynamics. However, in this instance, we are looking at a 200 time step subsection of the task. Hence, we can remove all superfluous network material which does not affect the behaviour and the dynamics of the network over these time steps using Algorithm 4. This achieves a significant reduction in network size, where the reduced network contains 7 genes and 5 epigenetic molecules. Using the tools developed during this work, we can query the network dynamics at specific time steps which can be seen in Table 2. These statistics show a stable topology leading up to the 1200th time step, where the decision to change the dynamics of the network is encountered. The network then partitions itself into two intermediary stable topologies and then onto a continuing stable topology for 138 time steps. We can see from the inputs in Table 2 that only inputs 1 and 2 are being used during this subsection of time steps.

The reason for the topological changes, and in turn a change in the output dynamics is the activation of gene 4, and both events are exactly correlated. What caused the activation of this gene? From analysis of the data of all the epigenetic molecules movements, 3 of the 5 are static meaning that they express the same genes regardless of internal or external perturbations. Out of the 2 dynamical epigenetic molecules, only one is within proximity of gene 4, and therefore is responsible for its expression. However, what caused




Figure 4: The decision of the optimum epiNet (0.8231) at each time step of the test data. Over layed is the optimum behaviour. The decision highlighted is made by epiNet to prematurely change its dynamics and incorrectly change its output from eyes open (1) to eyes closed (0).

Execution summary of epiNet on EEG classification					
Start	End	Length	Topology	Inputs	Output
1024	1199	175	0,1,3,5,6	1,2	1
1200	1219	19	0,1,3,4,5,6	1,2	0
1220	1225	5	0,1,2,3,4,5,6	1,2	0
1226	1364	138	0,1,3,4,5,6	1,2	0

Table 2: The summary of the execution of the decision of interest highlighted in Figure 5. It can be seen that only inputs 1 and 2 are being used throughout this period and that there are 4 topological changes of which one coincides with a change in output (decision).

the shift in this epigenetic molecule which resulted in the expression of gene 4? The epigenetic molecule was connected to, and tracking input 2 from the task. This can be seen in Figure 5.

We note that it's gene number 4 which is introduced, which ultimately causes the topological change. What caused the introduction of gene 4 is the movements of the 2nd epigenetic molecule which is transforming the input from input 2, to produce its movement, and introduce gene 4 (see Figure 5). This change which can be seen at time step 200 is significantly correlated between the epigenetic molecule and input 2.

Using this deductive reasoning, we can say that due to the fluctuation of input 2, which is connected to epigenetic molecules 2, this results in the changes to the network topology by expressing gene 4, which changes the dynamics of the network and ultimately the output of the network from 1 to 0.




Figure 5: The signal from input 2 over layed with the movements of the epigenetic molecule in which it is connected to. The topological changes occurs at approximately time step 1200, coinciding with a shift in position of the epigenetic molecule which the expresses gene 4. The input and epigenetic movement are highly correlated but not identical, as the epigenetic molecule is able to process the input according to its own regulatory function, alongside being connected to other genes.

Deducing Rules From The ‘Ultrasound Data of a Wall Following Robot’ Task

The decision that will be analysed for the robotic navigation task can be seen in Figure 6, where the optimum network was applied. This network achieved 100% on the test set, so all decisions it made coincide with the optimum behaviour highlighted in Algorithm 3. When the optimum epiNet is applied to steps 1-8 in the transparency algorithm, 14 genes are removed from the network, but all epigenetic molecules remain as they are essential to the dynamical behaviour of the network. The summary for this execution shows that only 2 inputs from the task are used to solve it optimally (Table 3), which coincides well with the optimal behaviour which is highlighted in Algorithm 3, where only the left and forward ultrasound distances are used.

The topological change that we are focusing on can be seen in Table 3 where gene 1 is suppressed from execution at time step 1150. Gene 4 is located within the region of the 2nd epigenetic molecule in the network’s movements throughout execution. No other epigenetic molecules interfere with this space, hence we can deduce that the 2nd epigenetic molecule is responsible for the suppression of gene 4.

In order to find out why the epigenetic molecule moved at the point it did, resulting in a decision change, the answer lies in which inputs from the task the epigenetic molecule is connected to. In this instance, the epigenetic molecule is connected to input 1, which is responsible for detecting the distance between the left hand side of the robot and a wall.

Execution summary of epiNet on robotic task					
Start	End	Length	Topology	Inputs	Output
1150	1210	60	0,1,2,3,4	0,1	0
1211	1250	49	0,2,3,4	0,1	1

Table 3: The data of the topological change responsible for the variation in dynamics seen in the decision of interest in Figure 7.




Figure 6: The output of the optimum epiNet applied to the wall following tasks, which 100% fitness was achieved. The highlighted region specifies a decision made by epiNet which will be analysed to discover why it occurred.

This connection is passed through a gene which effectively inverts the signal. This relationship can be seen in Figure 7, where only a small change in the input 2 signalling that a wall is getting closer, slightly adjusts the position of the epigenetic molecule which in turn moves it away from gene 4. This effectively shows that the topological change in the network was due to a change in the input, signalling that there is a wall on the left hand side. This topological change resulted in a change of dynamics as well as a change in output, instructing the robot to move to the right.

Although the information analysed in this section is a small element of the entire decision making process, it is possible this can be repeated for every decision made by the network. In this instance, as the network achieved the optimal behaviour, this will re-construct the original algorithm applied to the robot from Algorithm 3. In this sense epiNet twinned with thorough analysis is enough to build up a perfect picture of why the network behaves as it does over all time steps.

Conclusions

In this paper we investigate the benefits of epiNet, a machine learning architecture capable of self-topological modification. The topological modification of epiNet allows it to au-

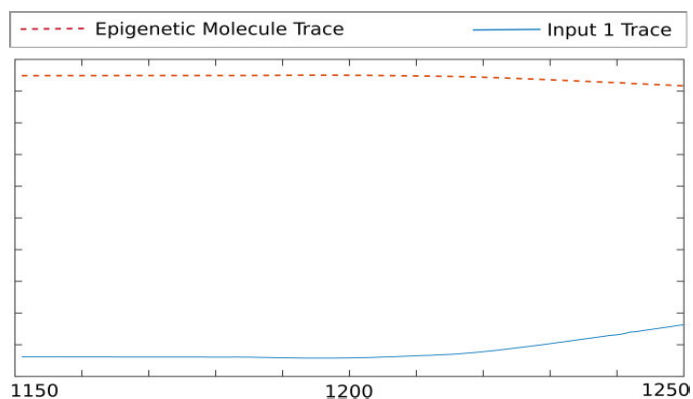


Figure 7: The correlation between the change in input and the inverse movement of the epigenetic molecule which is ultimately responsible for the decision made in figure 6.

tonomously decompose complex tasks, applying relatively small sections of its architecture to these subsections. This paper describes methods to exploit this behaviour to provide transparency of execution, to be able to question behaviours of the network and find out a logical answer as to why it behaves like it does. This work built a method to compile the data from the functionality of epiNet over varying time scales, by removing superfluous parts of the network which were not required to maintain functionality, whilst analysing the parts that are.

The work in this paper serves as a proof of concept of an architecture which is capable of autonomously revealing why it is making certain decisions and future work will follow two trends. Firstly how much of a complex problem can epiNet solve whilst providing transparency of its functionality? Secondly, how can we translate the work done in this paper to provide an autonomously generated, human readable description of network function? The results and analysis in this paper show that there is potential for answering both of these questions.

References

Capuano, F., Mulleter, M., Kok, R., Blom, H. J., and Ralser, M. (2014). Cytosine dna methylation is found in *drosophila melanogaster* but absent in *saccharomyces cerevisiae*, *schizosaccharomyces pombe*, and other yeast species. *Analytical chemistry*, 86(8):3697–3702.

Chambayil, B., Singla, R., and Jha, R. (2010). Virtual keyboard bci using eye blinks in eeg. In *2010 IEEE 6th International Conference on Wireless and Mobile Computing, Networking and Communications*, pages 466–470. IEEE.

Clark, D. and Pazdernik, N. (2013). *Molecular Biology*. Academic Press. Academic Press.

Emotive (2017). <https://www.emotiv.com/>.

Gentry, M. and Hennig, L. (2014). Remodelling chromatin to shape development of plants. *Experimental cell research*, 321(1):40–46.

He, K., Zhang, X., Ren, S., and Sun, J. (2015). Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034.

Lones, M. A., Turner, A. P., Fuente, L. A., Stepney, S., Caves, L. S., and Tyrrell, A. M. (2013). Biochemical connectionism. *Natural Computing*, 12(4):453–472.

Manshian, B. B., Pfeiffer, C., Pelaz, B., Heimerl, T., Gallego, M., Moller, M., Del Pino, P., Himmelreich, U., Parak, W. J., and Soenen, S. J. (2015). High-content imaging and gene expression approaches to unravel the effect of surface functionality on cellular interactions of silver nanoparticles. *ACS nano*, 9(10):10431–10444.

Nguyen, A., Yosinski, J., and Clune, J. (2015). Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 427–436. IEEE.

Parmeggiani, A., Barcia, G., Posar, A., Raimondi, E., Santucci, M., and Scaduto, M. C. (2010). Epilepsy and eeg paroxysmal abnormalities in autism spectrum disorders. *Brain and Development*, 32(9):783–789.

Smith, C. M., Finger, J. H., Hayamizu, T. F., McCright, I. J., Xu, J., Berghout, J., Campbell, J., Corbani, L. E., Forthofer, K. L., Frost, P. J., et al. (2014). The mouse gene expression database (gxd): 2014 update. *Nucleic acids research*, 42(D1):D818–D824.

Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. (2015). Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9.

Turner, A. P., Caves, L. S. D., Stepney, S., Tyrrell, A. M., and Lones, M. A. (2017). Artificial epigenetic networks: Automatic decomposition of dynamical control tasks using topological self-modification. *IEEE Transactions on Neural Networks and Learning Systems*, 28(1):218–230.

Turner, A. P., Lones, M. A., Fuente, L. A., Stepney, S., Caves, L. S., and Tyrrell, A. (2013a). The artificial epigenetic network. In *Evolvable Systems (ICES), 2013 IEEE International Conference on*, pages 66–72. IEEE.

Turner, A. P., Lones, M. A., Fuente, L. A., Stepney, S., Caves, L. S., and Tyrrell, A. M. (2013b). The incorporation of epigenetics in artificial gene regulatory networks. *BioSystems*, 112(2):56–62.

University of California, Irvine, School of Information and Computer Sciences (2017). Machine learning repository, <http://archive.ics.uci.edu/ml/datasets/>.

Wairagkar, M., Zoulias, I., Oguntosin, V., Hayashi, Y., and Nasuto, S. (2016). Movement intention based brain computer interface for virtual reality and soft robotics rehabilitation using novel autocorrelation analysis of eeg. In *Biomedical Robotics and Biomechatronics (BioRob), 2016 6th IEEE International Conference on*, pages 685–685. IEEE.

Zemach, A., McDaniel, I. E., Silva, P., and Zilberman, D. (2010). Genome-wide evolutionary analysis of eukaryotic dna methylation. *Science*, 328(5980):916–919.