

Information Visualization for Knowledge Extraction in Neural Networks

Liz Stuart¹, Davide Marocco², Angelo Cangelosi¹

¹ School of Computing Communication and Electronics
University of Plymouth, Drake Circus, PL4 8AA, Plymouth, UK
{lstuart, acangelosi}@plymouth.ac.uk

² Institute of Cognitive Science and Technologies
National Research Council, viale Marx 14, 00137, Rome, Italy
davide.marocco@istc.cnr.it

Abstract. In this paper, a user-centred innovative method of knowledge extraction in neural networks is described. This is based on information visualization techniques and tools for artificial and natural neural systems. Two case studies are presented. The first demonstrates the use of various information visualization methods for the identification of neuronal structure (e.g. groups of neurons that fire synchronously) in spiking neural networks. The second study applies similar techniques to the study of embodied cognitive robots in order to identify the complex organization of behaviour in the robot's neural controller.

1 Introduction

Knowledge extraction using neural networks typically involves the use of analytical methods for the automatic identification of information relevant to specific research goals. For example, the utilization of a finite union of open polytopes permits the transparent expression of knowledge embedded in recursive determinist perceptron [1]. In this paper, a complementary method for knowledge extraction from artificial and natural neural networks is presented. This is characterized by the active role of the researcher in the exploration of the neural network representation and the search for in-depth knowledge. This method is based on recent research and tools for information visualization in neural systems.

Information visualization [2], [3] is one of the fields of computer science that deals with the innovative representation of vast quantities of data. Consequently, it exploits advances in interactive computer graphics hardware, mass storage, and data visualization in order to visualize information. One of the fundamental principles of this field is the role of the investigator interacting with the data being analyzed and their ability to steer the exploration, in order to achieve greater insight. Thus, the investigator needs to be able to navigate throughout the whole dataset, in order to identify and explore specific subsets of interest. However, when visualizing large datasets, the issue of efficient navigation is amplified. It is important that the user is able to move to points of interest quickly without becoming disoriented within the dataset. Thus, it can be beneficial to restrict the mechanisms by which the user can

navigate within the dataset. Therefore, the user can be constrained to follow predetermined paths throughout the data space. In addition to navigation functionality, the investigator should also have control over the data representation itself. Thus, in order to truly steer the analysis, the investigator should be able to manage and overview the whole dataset, they should be able to filter and manipulate the data, select non-sequential subsets of interest and ultimately “drill-down” to inspect the actual data values that underpin the data represented. This follows the design principles of the much cited “information seeking mantra” coined by Shneiderman.

Recent research has focused on the development of new information visualization techniques and tools for neural networks. This includes the visualization of information from neuroscience research and the analysis on neural activity in embodied cognitive systems. In this paper, two such case studies are described. The first study demonstrates the use of information visualization methods for the identification of structure (e.g. groups of neurons that fire synchronously) in spiking neural networks. The second study applies the same techniques to embodied cognitive robots in order to identify the complex organization of sensorimotor behaviour and its management by a neural controller.

2 VISA: Information Visualization for Spiking Neural Networks

The key to numerous issues within the field of neuroscience is linked with the theoretical understanding of vast quantities of experimental neural data. In particular, investigation of information processing in the nervous system is associated with the analysis of this vast resource of neural data, namely, simultaneously recorded multi-dimensional spike train data.

Much of the research focus in this area is focused around the principle of synchronization of neural activity [4], [5]. However, the experimental evidence that is currently available requires further, in-depth analysis in order to extract the knowledge inherent in these datasets. It is clear that analysis of neural data such as multi-dimensional spike trains using traditional tools like raster plots and cross-correlograms increases in complexity as the size of the datasets increase. Therefore, new methods of analyzing this data, designed specifically for large datasets, are required.

Traditionally, analysis of multi-dimensional spike train data has not supported real-time user interaction. In 1996, Shneiderman [2] identified user interaction as a primary essential component of information visualization representations. Shneiderman also introduced the “information-seeking mantra” that highlighted user requirements in this area. It specified that users should have the capability to overview data in order to see the whole dataset in a single display. The mantra also recommended that users needed to be able to zoom in on “interesting” areas of their datasets and to filter out parts of the datasets not required for the current investigation. Finally, the mantra specified that users needed to be able to get details-on-demand, access to the fundamental data items which were used to create the visual representation.

Based on this mantra, a toolbox of interactive methods for exploring neural data was developed by Stuart and collaborators [6] as part a research project called VISA, Visualization of Inter-Spike Associations. This toolbox facilitates zooming, filtering and manipulation of neural data and supports the use of multiple views of the same data, as well as real time interaction. One of the key visualization tools currently included within VISA is subsequently described.

2.1 Traditional parallel coordinates

Interest in parallel coordinates was rejuvenated by Inselberg et al. [7]. The true value of parallel co-ordinates is their ability to represent vast quantities of multi-variate data in a simply 2-d representation. Traditional presentation of parallel coordinates denotes a series of data points as vertical axis coordinate values distributed along a horizontal axis.




Figure 1 Representation of the points $p(2,1,3)$ and $q(3,2,0)$ using parallel coordinates.

Thus, a specific point in n -dimensional Euclidean space is represented by n vertical axes values distributed along the horizontal axis. To illustrate this, consider the two 3-dimensional points $p(2,1,3)$ and $q(3,2,0)$. Refer to Figure 1 for an illustration of these two points in 3-dimensional space represented in parallel coordinates.

2.2 The use of parallel coordinates in VISA

Parallel coordinates are traditionally used to identify correlations between variables and to convey aggregation information. Although the VISA parallel coordinate tool [8] is based on the original method, it has been adapted for use in the analysis of spiking activity in neural networks.

In brief, the original multi-dimensional spike train dataset is transformed using a well known and respected method of analysis called the gravity transform algorithm defined by Gerstein et al [9]. Thus, the data under investigation is now the position of n particles in n -dimensional space where n is the number of neurons initially specified. The traditional means of displaying output from the Gravity Transform does not scale up easily for large numbers of particles. Thus, parallel coordinates were introduced to represent the n -dimensional positions of the n particles. Therefore, all n particle parallel coordinates were displayed simultaneously on the display to represent the position of all n particles at a specific point in time. Subsequently, animation was used to represent the movement of the particles over time. Numerous trials have been carried out using this technique to support the identification of

neuronal assemblies from multi-dimensional spike train datasets. The success of these trials is well documented [8,10]. Furthermore, this work continues as the parallel coordinate tool undergoes further adaptation [11].

3 Sensorimotor Knowledge Integration in the Neural Controller of Cognitive Robots

Research in cognitive systems, including natural (animals and humans) and artificial (agents, robots) systems, supports increased understanding of the relationship between cognitive, neural, social and evolutionary factors. Various researchers are working on the design of cognitive robots that have sensorimotor capabilities, to interact with their environment, and cognitive and linguistic skills to build internal representations of their physical social environment and talk about them [12]. The complex patterns of interaction between the different sensorimotor and cognitive capabilities require the development of new methodologies to increase understanding of the cognitive systems. Information visualization has a role to play in the analysis of these complex sets of behavioural and cognitive data [13].

A new study based on the application of parallel coordinates to the visualization of activity in the neural controllers of linguistic cognitive robots [14] is presented. The primary aim is to investigate the interaction between sensory neurons and the internal categorical and linguistic representations in simulated robotic agents that are able to interact with two different objects (i.e. touch a sphere and avoid a cube). These agents are also able to develop, through evolution, a shared lexicon to name the two objects.

The cognitive robotic model consists of a 3-segment arm with 6 degrees of freedom (DOF). Each segment consists of a basic structure of two cylindrical bodies and two joints. This structure is replicated three times, with the final segment being shorter to represent the hand of the robot (with no fingers). The controller of each individual robot consists of an artificial neural network with 11 sensory neurons connected to 3 hidden neurons. These connect to 8 output neurons. The first 9 sensory neurons encode the angular position (normalized between 0.0 and 1.0) of the 6 DOFs and the state of three contact sensors located in the three corresponding arm segments. The other 2 sensory neurons receive their input from the other agents (name of objects). The first 6 motor neurons control the actuators of the corresponding joints. The motor is activated so that it is able to apply a force proportional to the difference between the current and the desired position of the joint. The last 2 output neurons encode the signal to be communicated to the other agents. Agents are evolved for their ability to interact with the objects, using a genetic algorithm.

To investigate the neural control strategies for the robot's sensorimotor and cognitive behaviour, a dataset from two experimental trials was created. In the first trial, the activity of all the input, hidden and output units were recorded whilst the robot interacted with a sphere. The second trial records the neural activity during the interaction of the robot with the cube. Each trial produced a dataset of 16 variables consisting of the 4 linguistic units, the 9 input units (3 touch sensors and 6 proprioceptive), and the 3 hidden nodes. In addition, each trial lasted for 150 cycles of actions (i.e. neural networks activations), thus producing 150 x 16 data points. All

activation values (2 x 16 x 150) were plotted in the parallel coordinate representation. This was produced by a modified version of the VISA software.

This parallel coordinate tool produces a dynamic and interactive representation of the whole dataset. For example, it is possible to search for the different patterns of activation that distinguish the two behavioural tasks. Initially, the whole dataset, 150 x 16 cycles, was plotted using a different colour to represent each of the two tasks. During experimentation, the activation cycles were identified and eliminated. Note: these were the lines (parallel coordinates) from the two tasks which were coincident (i.e. the same unit has the same activation value in both tasks). This process of eliminating redundant (less informative) cycles gradually revealed the small number of critical cycles in the display. These are the lines (parallel coordinates) from each task that have distinct activation patterns.




Figure 2: Parallel coordinate display for cognitive robotics. See text for explanation.

Figure 2 shows a snapshot of the dynamic information visualization process with the data from the robotic experiment. The display shown solely depicts the six critical cycles of interaction (cycles 65-70 for each interaction) when the robot arm makes contact with the two objects (light gray for spheres, dark for cubes). This display clearly shows, with a single view, which input and hidden units are involved values in the two tasks. During interaction with the sphere, the units active solely for the touching behaviour are the first 2 proprioceptive sensors (vertical axes 5 and 6), the touch sensor of the second arm segment (axis 10) and the 3rd hidden unit (the final vertical axes). Instead, the units that specialize for the cube avoiding behaviour are the pairs of proprioceptive sensors for the 2nd and 3rd segment (respectively units 8-9 and 11-12) and the 1st hidden units (vertical axis 14).

4 Conclusions

The two studies presented in this paper demonstrate the usefulness of information visualization methods for knowledge extraction in various types of neural network. In the first study, the parallel coordinate method included in the VISA software project has supported the identification of common activity in networks of spiking neurons. In the second study, research on cognitive robots and the further adapted parallel

coordinate tool was used to identify the input and output units involved in the neural control of the sensorimotor behaviour in robots.

The work presented here is an innovative approach to knowledge extraction in artificial neural networks. Instead of relying solely on formal methods for knowledge elicitation, this approach is based on an active exploration and visualization of neural network data. This design is largely based on Shneiderman's information-seeking mantra, where users have the capability to "overview data"; to "zoom and filter" data and also to obtain "details-on-demand".

Future research is looking at the development of new information visualization tools for neural networks and cognitive robotics research as well as the further use of current tools. For example, extension of the parallel coordinate tool is required in order to support greater investigation of the various interrelationships between the neural network units. To achieve this, the tool will require the additional functionality to enable the user to move variables (the vertical axes) along the horizontal axis. This will enable users to analyse the precise relationships between any two variables which are not adjacent by default in the original display. In addition, the user interface will be further developed. It is important that the user is able to specify the colours used to represent both the lines (parallel coordinates) corresponding to correct and incorrect behaviours, and to highlight differences in the pattern of the neural network activation.

References

- [1] Tajine, M. & Elizondo, D. (1998), Growing methods for constructing Recursive Deterministic Perceptron neural networks and knowledge extraction. *Artificial Intelligence*, 102, 295-322.
- [2] Shneiderman, B. (1996), The eyes have it: A task by data type taxonomy of information visualizations, *Proc. IEEE Symposium on Visual Languages '96*, IEEE, 336-343.
- [3] Spence R. Information Visualization. Addison-Wesley: Harlow, UK, 2001.
- [4] Borisyuk RM & Borisyuk GN (1997), Information coding on the basis of synchronisation of neuronal activity, *BioSystems*, 40, 3-10.
- [5] Fries P, Neuenschwander S, et al. (2001), Rapid feature selective neuronal synchronization through correlated latency shifting, *Nature Neuroscience*, 4(2), 194-200.
- [6] <http://www.plymouth.ac.uk/infovis>
- [7] Inselberg A & Dimsdale B (1990), Parallel Coordinates: A tool for visualising multidimensional geometry, *Proc. Visualization'90*, 361-378.
- [8] Stuart L., Walter M. and Borisyuk R. (2002), Visualization of synchronous firing in multi-dimensional spike trains, *BioSystems* 67:265-279.
- [9] Gerstein G.L. & Aertsen A.M. (1985), Representation of cooperative firing activity among simultaneously recorded neurons, *Journal of Neurophysiology*, 54(6), 1513-1528.
- [10] Stuart, L., Walter, M. and R. Borisyuk (2001), Visualization of multi-dimensional Spike Trains, *Proc. 4th International workshop Neural Coding'2001*, 47-48.
- [11] Barlow, N. & Stuart, L. (2004), Animator: A Tool for the Animation of Parallel Coordinates, *Proc. IEEE Intl. Conference on Information Visualization*, IV04, 725-730.
- [12] Cangelosi, A. & Parisi, D., (Eds.) (2002), *Simulating the Evolution of Language*, Springer.
- [13] Smith, T., Bullock, S. & Bird, J. (2002), Beyond fitness: Visualising evolution - Workshop overview, *ALife VIII*, Sydney.

- [14] Marocco, D., Cangelosi, A. & Nolfi S. (2003), The emergence of communication in evolutionary robots. *Philosophical Transactions of the Royal Society of London – A*, 361, 2397-2421.