# A Walsh Analysis of Multilayer Perceptron Function

Kevin Swingler

*Computing Science and Maths, University of Stirling, Stirling, FK9 4LA, Scotland*

Keywords: Multilayer Perceptrons, Walsh Functions, Network Function Analysis.

Abstract: The multilayer perceptron (MLP) is a widely used neural network architecture, but it suffers from the fact that its knowledge representation is not readily interpreted. Hidden neurons take the role of feature detectors, but the popular learning algorithms (back propagation of error, for example) coupled with random starting weights mean that the function implemented by a trained MLP can be difficult to analyse. This paper proposes a method for understanding the structure of the function learned by MLPs that model functions of the class $f : \{-1, 1\}^n \to \mathbb{R}^m$. The approach characterises a given MLP using Walsh functions, which make the interactions among subsets of variables explicit. Demonstrations of this analysis used to monitor complexity during learning, understand function structure and measure the generalisation ability of trained networks are presented.

## 1 INTRODUCTION

The multilayer perceptron (MLP) (Rumelhart et al., 1986) is a widely used neural network architecture. It has been applied to regression, classification and novelty detection problems and has been extended in various ways to process time varying data, e.g. (Elman, 1990). In the field of data mining MLPs are a common choice amongst other candidates such as classification trees, support vector machines and multiple regression. Due to the wide variety of tasks for which they are suited, and their ability as universal approximators, MLPs have become very popular. However, there is one aspect of the MLP that restricts and complicates its application, and that is the role of the hidden neurons. A common criticism of the MLP is that its knowledge is not represented in a human readable form. The comparison that is often made is with classification or regression trees, which represent partitions in the input space explicitly in their structure. This makes human understanding of the underlying function and the reasons behind any given output quite easy. Given a picture of a classification tree, a human may apply it to an input pattern without even needing a computer to run the algorithm. This is far from simple with an MLP.

The hidden units in an MLP act as feature detectors, combining inputs from below into higher order features that are, in turn, combined by higher layers still. The common learning algorithms such as back propagation of error (Rumelhart et al., 1986) have no explicit means of ensuring that the features are optimally arranged. Different neurons can share the same feature, or have overlapping representations. In networks where each layer is fully connected to the one above, every hidden neuron in a layer shares the same receptive field, so their roles often overlap. This makes analysis even more difficult as hidden neurons do not have independent roles. The inclusion of additional layers of hidden neurons compounds the problem further.

Some work has been carried out on the analysis of hidden neurons in MLPs. For example, (Kamimura, 1993) used an entropy based analysis to identify important hidden units (known as principal hidden units) in a network for the purpose of pruning an oversize hidden layer. (Sanger, 1989) proposed a method of contribution analysis based on the products of hidden unit activations and weights and (Gorman and Sejnowski, 1988) presented a specific analysis of the hidden units of a network trained to classify sonar targets.

The question of how to extract rules from multilayer perceptrons has received more attention and is still a very active field of research. (Kulluk et al., 2013) propose a fuzzy rule extraction method for neural networks, which they call Fuzzy DIFACONN. (Hruschka and Ebecken, 2006) propose a clustering based approach to MLP rule extraction that uses genetic algorithm based clustering to identify clusters of hidden unit activations which are then used to generate classification rules. (Saad and Wunsch II, 2007) use an inversion method to generate rules in the form of hyperplanes. Inverting an MLP (i.e. finding the inputs that lead to a desired output) is done by gradient descent and using an evolutionary algorithm.

Both (Augasta and Kathirvalavakumar, 2012) and (Jivani et al., 2014) present recent comparative studies of neural network rule extraction, distinguishing between methods that are decompositional, pedagogical and eclectic. A decompositional approach extracts rules from the weights and activations of the neural network itself. The pedagogical approach, which is taken in this paper, treats the neural network as a black box and generates rules based on the outputs generated by the network in response to a set of input patterns. Eclectic rule extraction combines both of the aforementioned approaches.

More work has concentrated on choosing the right number of hidden units for a specific data set. (Baum and Haussler, 1989) bound the number of weights by the target error size, (Uphadyaya and Eryurek, 1992) bounded the number of hidden units by the number of patterns to be learned, (Widrow and Lehr, 1990) chose a bound based on the number of output units in the network, and (Weigend et al., 1992) pointed out that the amount of noise in the training data has an impact on the number of units used. Some have taken a dynamic approach to network structure discovery, for example (Bartlett, 1994) used an information theoretic approach to add or remove hidden neurons during training. The problem with this approach to training an MLP is that the existing weights are found in an attempt to minimise error for that number of hidden neurons. Adding a new one may mean the existing weights are starting in a configuration that is unsuitable for a network with more neurons. Other search methods have also been applied to finding the right structure in an MLP. (Castillo et al., 2000) and (Yao, 1999) used genetic algorithms to search the space of network structures, for example.

When using MLPs (and other machine learning techniques), it is common practice to produce several models to be used in an ensemble (Krogh and Vedelsby, 1994). Due to the random start point of the weight values, and the differences in architecture across the networks in an ensemble, it is not easy to know whether or not different networks are functionally different. It is possible to train a number of different MLPs that all implement the same function (perhaps with differing quality of fit across the weights) with very different configurations of weight values. For example, one could re-order the hidden units of any trained network (along with their weights) and produce many different looking networks, all with identical functionality. One way to compare MLPs is to compare their outputs, but a structural comparison might also be desirable, and that is what we present here.

Note the distinction between the structure of an MLP, which is defined by the neurons and connecting weights, and the structure of the function it implements, which can be viewed in a number of other ways. This paper views the underlying function implemented by an MLP in terms of the contribution of subsets of input variables. The number of variables in a subset is called its order, and there are $\binom{n}{k}$ subsets of order $k$ in a network of $n$ inputs. The first order subsets are the single input variables alone. The second order subsets are each of the possible pairs of variables, and so on. There is a single order $n$ set, which is the entire set of inputs. Any function can be represented as a weighted sum of the values in each of these subsets. The weights (known as coefficients in the chosen analysis) are independent (unlike the weights in an MLP, whose values are determined to an extent by other weights in the network) and specific to their variable subset. The first order coefficients describe the effect of each variable in isolation, the second order coefficients describe the contribution of variable pairs, and so on. The method for decomposing a neural network function into separate components described in this paper is the Walsh transform.

Section 2 describes the Walsh transform in some detail. This is followed by a description of the method for producing Walsh coefficients from a neural network in section 3. Section 4 introduces some functions that will be used in experiments described in following sections. Section 5 demonstrates how the method can be used to track the complexity of MLPs during training and section 6 demonstrates how a partial transform on a small sample from a larger network can provide useful insights. The Walsh method is compared to other methods of understanding network structure in section 7. Finally, sections 8 and 9 offer some conclusions and ideas for further work.

## 2 WALSH FUNCTIONS

Walsh functions (Walsh, 1923), (Beauchamp, 1984) form a basis for real valued functions of binary vectors. Any function $f : \{-1,1\}^n \rightarrow \mathbb{R}$ can be represented as a weighted linear sum of Walsh functions. The Walsh functions take the form of a sequence of bit strings over $\{-1,1\}^{2^n}$ where $n$ is the number of variables in the function input. $n$ is known as the Walsh function order. There are $2^n$ Walsh functions of order $n$, each $2^n$ bits long. Figure 1 shows a representation of the order 3 Walsh functions. Each Walsh function has an index from 1 to $2^n$, with the $j^{\text{th}}$ function being $\psi_j$ and bit number $x$ of the $j^{\text{th}}$ Walsh function is $\psi_j(x)$. As figure 1 shows, the Walsh functions can be viewed as a matrix of values from $\{-1,1\}$ with rows
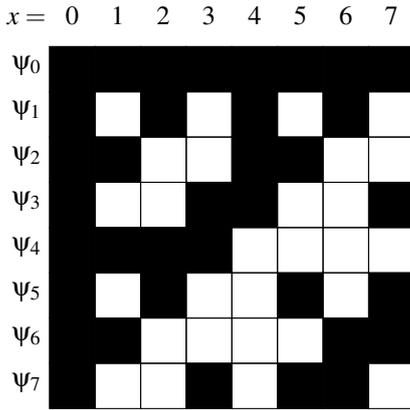
$x = $ 0 1 2 3 4 5 6 7



Figure 1: A pictorial representation of an order 3 Walsh matrix with black squares representing 1 and white squares -1. A Walsh sum is calculated by summing the product of the Walsh coefficient associated with each row by the values in the column indexed by the function input.

representing each Walsh function and columns representing each bit.

A Walsh representation of a function $f(\mathbf{x})$ is defined by a vector of parameters, the Walsh coefficients, $\omega = \omega_0 \ldots \omega_{2^n-1}$. Each $\omega_j$ is associated with the Walsh function $\psi_j$, that is a row in the Walsh matrix. Each possible input, $\mathbf{x}$ is given an index, $x$, which is calculated by replacing any -1 in $\mathbf{x}$ with 0 and converting the result to base 10. For example if $\mathbf{x} = (1,-1,1)$, then $x = 5$. Each column of the Walsh matrix corresponds to a value of $x$.

The Walsh representation of $f(\mathbf{x})$ is constructed as a sum over all $\omega_j$. In the sum, each $\omega_j$ is either added to or subtracted from the total, depending on the value of the bit corresponding to $x$ (i.e. column $x$ in the Walsh matrix), which gives the function for the Walsh sum:

$$f(\mathbf{x}) = \sum_{j=0}^{2^n} \omega_j \psi_j(x) \qquad (1)$$

### 2.0.1 Constructing the Walsh Functions

The value of a single cell in the Walsh matrix, $\psi_j(x)$ is calculated from the binary representation of the coordinates $(j,x)$, of $\mathbf{j}$ and $\mathbf{x}$, and returns +1 or -1 depending on the parity of the number of 1 bits in shared positions. Using logical notation, a Walsh function is derived from the result of an XOR (parity count) of an AND (agreement of bits with a value of 1):

$$\psi_j(x) = \oplus_{i=1}^{n} (x_i \wedge j_i) \qquad (2)$$

where $\oplus$ is a parity operator, which returns 1 if the argument list contains an even number of 1s and -1 otherwise.

### 2.0.2 Calculating the Coefficients - the Walsh Transform

The Walsh transform of an $n$-bit function, $f(\mathbf{x})$, produces $2^n$ Walsh coefficients, $\omega_x$, indexed by the $2^n$ combinations across $f(\mathbf{x})$. Each Walsh coefficient, $\omega_x$ is calculated by

$$\omega_x = \frac{1}{2^n} \sum_{j=0}^{2^n-1} f(j)\psi_j(x) \qquad (3)$$

Each of the resulting Walsh coefficients has an index, which defines the set of input variables over which it operates. Converting the index to a binary representation over $n$ bits produces a representation of the variables associated with the coefficient where a 1 in position $i$ indicates that $x_i$ contributes to the effect of that coefficient. For example, over 4 bits, the coefficient $\omega_9$ produces a binary word 1001, which tells us that $x_1$ and $x_4$ contribute to the effect of $\omega_9$. The order of a coefficient is defined as the number of bits it contains that are set to 1. For example, $\omega_2$ and $\omega_8$ are first order as they have one bit set to 1, and $\omega_9$ is second order. The magnitude of a coefficient indicates its importance in contributing to the output of the function on average across all possible input patterns.

A function of $n$ inputs produces $2^n$ Walsh coefficients, so it is not always possible to consider the value of each coefficient individually. In this work we look at individual coefficients and also define some simple aggregate measures for summarising the results of a Walsh transform. They are the number of non-zero coefficients, which is taken as a crude measure of overall complexity, and the average magnitude of coefficients at each order, which produces a set of values that measure the contribution to the models output made on average by interactions of each possible order.

## 3 METHOD

In this context, the Walsh transform is not used to understand the training data, but to understand a neural network that was trained on that data. The analysis is in terms of the inputs to and the outputs from the network, not its weights or activations, making this a pedagogical approach. The black box of the neural function is assessed in terms of its Walsh decomposition. Walsh functions map a vector of binary valued inputs onto a real valued output, so any function with this structure is amenable to the analysis. As shown below, multiple output neurons and classification net-

works may also be analysed with this approach, so the outputs can be nominal, discrete or continuous.

As neural networks can generalise and produce an output for any given input pattern, we can generate an exhaustive or randomly sampled data set from which to perform the Walsh transform. A full Walsh decomposition, as defined in equation 3 requires an exhaustive sample of the input space. In all but the smallest of networks, this is unfeasible in an acceptable time period, so the coefficients must be calculated from a sample. In either case, the sample used to calculate the coefficients is generated from the whole input space, not just the training data. The significant coefficients (those that are significantly far from zero) can be very informative about the underlying structure of the function (in this case, the MLP). The procedure is similar to that of pedagogical rule discovery in that it treats the MLP as a black box and performs an analysis on the output values that the network produces in response to input patterns. The method proceeds as follows:

1. Build a single MLP using your chosen method of design and weight learning;

2. Generate input patterns (either exhaustively or at random) and allow the MLP to generate its associated output, thus producing $(\mathbf{x}, f(\mathbf{x}))$ pairs;

3. Use the resulting $(\mathbf{x}, f(\mathbf{x}))$ pairs to perform a Walsh transform using equation 3;

4. Analyse the significant coefficient values, $\omega_x$.

The method can also be used for MLPs designed for classification rather than regression. In such cases, there is normally a single output neuron for each class, with a target output value of one when the input belongs to the neuron's designated class and zero otherwise. Properly trained, each neuron represents the probability of a new pattern belonging to its designated class. Such a network is effectively a number of related functions (one for each class) with a continuous output between zero and one. Each output neuron can be analysed in turn using the same procedure.

Step 4, the analysis of the $\omega_x$ values can take many forms. This paper discriminates between analysis during training (section 5) where the goal is to gain an insight into the level of complexity a network achieves as learning progresses, and post training analysis, designed to provide insights into the function of the trained network. The example of such an analysis in section 6.1 shows how the generalisation ability of a network may be investigated from the results of the Walsh analysis. The goal of the analysis is not to generate rules, so this is not another rule extraction method, rather it is designed to give human insights into the hidden life of the MLP.

## 4 EXPERIMENTS

A set of functions of increasing complexity[1] were chosen to generate data to test this analysis. They are:

**OneMax**, which simply counts the number of values set to one across the inputs. This is a first order function as each variable contributes to the output independently of any others. The OneMax function is calculated as

$$f(\mathbf{x}) = \sum_{i=1}^{n} x_i \qquad (4)$$

**Vertical symmetry**, which arranges the bits in the input pattern in a square and measures symmetry across the vertical centre line. This is a second order function and is calculated as

$$f(\mathbf{x}) = \sum_{i=1}^{n} \sum_{j=1}^{n} \delta_{ij} s_{ij} \qquad (5)$$

where $\delta_{ij}$ is the Kronecker delta between $x_i$ and $x_j$, and $s_{ij}$ is 1 when $i$ and $j$ are in symmetrical positions and 0 otherwise.

**K-bit trap** functions are defined by the number of inputs with a value of zero. The output is highest when all the inputs are set to one, but when at least one input has a value of zero, the output is equal to one less than the number of inputs with a value of 0. A k-bit trap function over $n$ inputs, where $k$ is a factor of $n$ is defined by concatenating subsets of $k$ inputs $n/k$ times. Let $\mathbf{b} \in \mathbf{x}$ be one such subset and $c_0(\mathbf{b})$ be the number of bits in $\mathbf{b}$ set to zero.

$$f(\mathbf{x}) = \sum_{\mathbf{b} \in \mathbf{x}} f(\mathbf{b}) \qquad (6)$$

where

$$f(\mathbf{b}) = \begin{cases} c_0(\mathbf{b}) - 1, & \text{if } c_0(\mathbf{b}) > 0 \\ k, & \text{if } c_0(\mathbf{b}) = 0 \end{cases} \qquad (7)$$

The first case in equation 7, which applies to all but 1 in $2^k$ patterns, could be modelled with a first order network (a linear perceptron, for example), which is a local minimum in the error space. The 'trap' part of the function is caused by the second case in equation 7, which requires the output to be high when all of the inputs have a value of zero. This requires a higher order function, including components at orders from 1 to $k$, but only a small proportion of the data (1 in $2^k$ of them) contains any clue to this.

---

[1]Complexity has a specific meaning in this context. It describes the number and order of the interactions between inputs that produce a function's output.

# 5 ANALYSIS DURING TRAINING

Experiments were conducted to investigate the structure of the function represented by an MLP as it learns. The MLP used in these experiments had a single hidden layer and one linear output neuron. The functions described above were used to generate training data, which was used to train a standard MLP using the error back propagation algorithm. At the end of each epoch (a single full pass through the training data), a Walsh transform was performed on the predictions made by the network in its current state.

Summary statistics designed to reflect the complexity of the function the network has implemented and the level of contribution from each order of interaction were calculated from the Walsh coefficients. The complexity of the function was calculated as the number of significant non-zero Walsh coefficients. The size of the contribution from an order of interaction, $o$ was calculated as the average of the absolute value of the coefficients of order $o$. Experiment 1 trained networks on the simple OneMax function (equation 4). Figure 2 shows the training error and network complexity of an MLP with one hidden unit trained on the OneMax function. During learning, the network initially becomes over complex and then, as the error drops, the network complexity also drops to the correct level.

In experiment 2, an MLP was trained on the symmetry function of equation 5, which contains only second order features. Figure 3 shows the results of analysing the Walsh coefficients of the network function during learning. Three lines are shown. The solid line shows the network prediction error over time and the broken lines show the contribution of the first and second order coefficients in the Walsh analysis
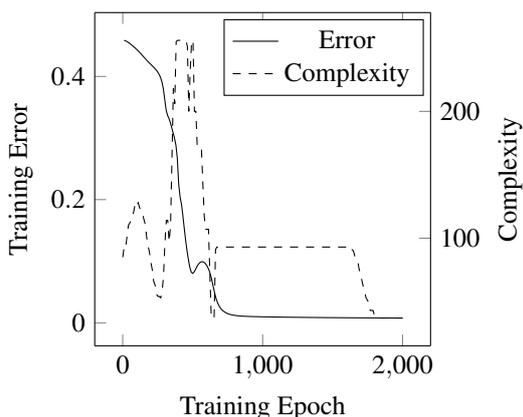


Figure 2: Comparing training error with network complexity during learning of the OneMax function with an MLP with one hidden unit. Note that complexity falls almost 1000 epochs after the training error has settled at its minimum.
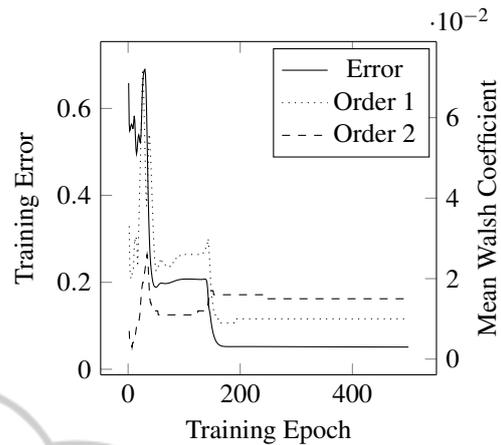


Figure 3: Network error and contribution of first and second order Walsh coefficients during training of an MLP on a second order function. Note the fall in the error rate when the second order coefficients overtake the first.

of the network function. Note the point in the error plot where the error falls quickly corresponds to the point in the Walsh analysis where the second order coefficients grow past those of first order. Compare this chart to that in figure 4, where the same problem is given to another MLP with the same structure, but which becomes trapped at a local error minimum, which is a first order dominated approximation to the function. The plot suggests that the higher order components cannot increase their contribution and that this network is unlikely to improve.

Figure 5 shows the error of an MLP decrease as it learns the 4 bit trap problem described by equation 6. The contribution of the first, second, third and fourth order Walsh coefficients are summed and plotted. The final, correct configuration can be seen in the right hand part of the plot, with the first order coefficients having the strongest contribution, but with the second, third and fourth also required to escape the 'trap' of the order below. The plot shows the first order coefficients growing first (as they did in figure 3), causing the average error to rise due to the higher order trap part of the function. The first order components are suppressed by the high error they cause, but the error doesn't settle to its lowest point until the first order coefficients recover the correct level of contribution.

# 6 PARTIAL WALSH ANALYSIS

For even moderately large numbers of inputs, calculating every Walsh coefficient can take an impractically long amount of time. In such cases, a partial Walsh analysis may still be useful. A partial analy-
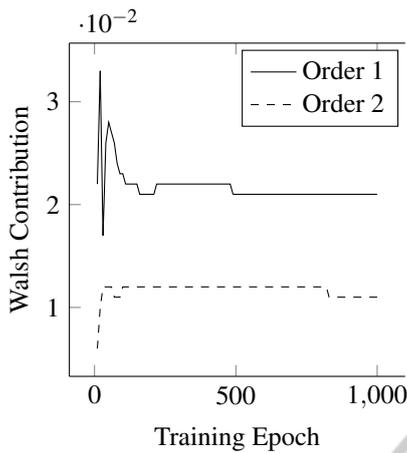
Figure 4: Contribution of first and second order Walsh coefficients during training of an MLP on a second order function, stuck in a local minimum.
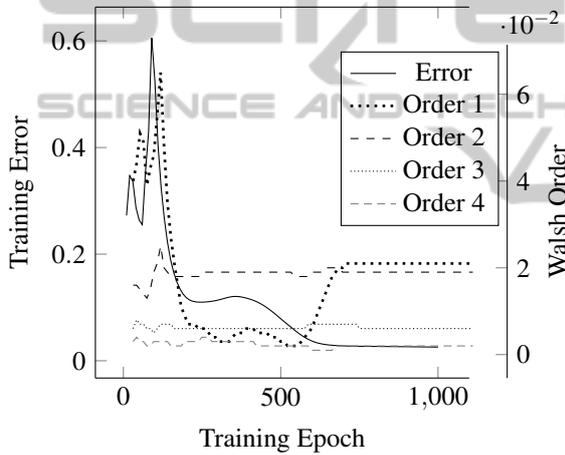


Figure 5: The contribution of first, second, third and fourth order Walsh coefficients during training of an MLP on a the 4-bit trap function plotted with the average error per pass through the data set (solid line).

sis calculates the values of only a small subset of the Walsh coefficients. An obvious choice for the subset of coefficients to calculate are those of the lower orders. $\omega_0$ is the average output of the function (in this case, the MLP) across the sampled data. The first order coefficients, $\omega_1, \omega_2, \omega_4, \ldots$ represent the average contribution of each input in isolation. In general, order $k$ coefficients represent the additional contribution of each subset of inputs of size $k$ to the function output. The number of coefficients of order $k$ from a set of inputs of size $n$ is $\binom{n}{k}$.

It is also possible to estimate the Walsh coefficients from a sample of random input patterns and their associated predicted outputs from the network, rather than analysing every input pattern exhaustively. As with the calculation of any statistic from a sam-

ple, the values gained are estimates, but they can still provide useful insights into the functioning of a neural network. The number of samples required to estimate coefficients accurately grows exponentially with their order, so the low order coefficients can be estimated with smaller samples than the higher order coefficients require. The next experiment described in this paper makes use of partial samples from both the coefficients and the input space.

## 6.1 Measuring Generalisation

The ability of an MLP to generalise to produce outputs for patterns that were not in its training data is of great advantage. As the weights of the network are difficult to analyse, the performance of the learned function in areas of input space that are outside those covered by the training data can be difficult to assess. Test and validation sets perform this task to a degree, but this paper proposes a new method based on a Walsh analysis.

The Walsh coefficients of an MLP function are generated by sampling (either exhaustively or at random) from the whole input space, not just the part of it covered by the training or test data. The coefficients give a picture of the general shape of the function, not just its behaviour on the training data. For ease of visualisation, a character classification task was used to test this type of analysis. Figure 6 shows the digits from 0 to 9 as 25 pixel bitmaps. A standard MLP with 25 binary inputs and 10 binary outputs was used to learn the classification. Data was generated using the images in figure 6 with evenly distributed random noise (bit values flipped at random) added at varying levels. The resulting MLP implements ten different functions that map the binary input patterns onto single, continuous outputs. With a suitable training regime the output values represent the probability of an input pattern belonging to the class represented by the output neuron. These functions are not independent – they share the parameters from the inputs to the hidden layer, but differ in their hidden to output weights. Ten Walsh decompositions are performed – one for each output neuron – and used to separate out the ten functions that are combined across the MLP's weights.

An individual Walsh decomposition for each output neuron was performed after training was complete based on 50,000 random input samples and their associated network output. The first order coefficients were plotted on a grid where the pixel locations from the inputs correspond to the first order coefficients of the Walsh decomposition.

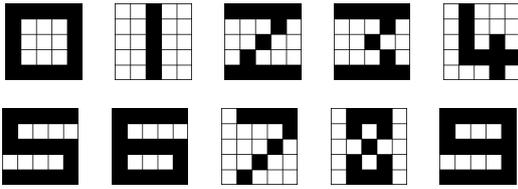Figure 7 shows the results for the network trained

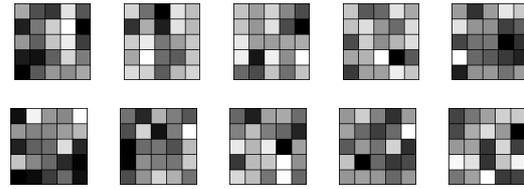Figure 6: Noise free training data used to analyse a classification network.



Figure 7: First order Walsh coefficients from a network trained on the data in figure 6 with no added noise. Grey squares indicate no contribution to classification from a first order component. Greater depth of black or white indicates stronger contribution (positive or negative).
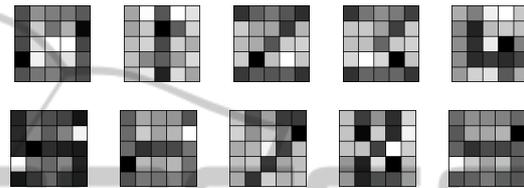


Figure 8: First order Walsh coefficients from a network trained on the data in figure 6 with 30 percent added noise. The noise ensures that no individual input can be relied upon to produce a correct classification, and so produces a model that covers more of the input space, and so is better at generalisation.

on noise free data. The network was able to separate the training data perfectly, but the first order coefficients of its underlying function do not suggest that a particularly general model has been learned. As the MLP needs only to find a weight configuration that minimises error, it will do so using a minimal set of features if possible. These features are difficult to identify by an analysis of the network weights, but the Walsh analysis reveals more. From figure 7 it is clear that a small number of pixels have been identified as key discriminators. These are the black squares in each image. For example, it is enough to know that there is a black pixel in the second row of the centre column to classify the image as a 1 (no other digit contains a black pixel there) and the coefficients for the output neuron of class '1' reflect that. Such a network is unlikely to generalise well as noise in the key inputs would lead to a misclassification.

Contrast this with the first order coefficients shown in figure 8, which were calculated from a network trained with 30% added noise. The patterns captured by the first order components of the network function mapped to each output are clearly more generally matched to the patterns they have been trained to identify. Figure 8 clearly shows that more of the inputs have a role in the classification function, suggesting a more robust model. Note the white squares in the images for patterns 5,6 and 9. They show the locations of key negative evidence for those classes. A value of one in these inputs is strong evidence against the pattern belonging to that class. These pixels are at the key points of difference between the patterns. A recent paper, (Szegedy et al., 2013) has found that deep neural networks have what the authors call "blind spots", which manifest themselves as images that are clearly in one class, but are erroneously classified by the network. Such examples are discovered using a method that starts with a correctly classified image and searches for the closest image that causes a mis-classification, known as an adversarial image. It is clear from figure 8 that the Walsh analysis can highlight the key pixels which, if changed, will produce an adversarial image. This work suggests (but leaves for future investigation) that the blind spots could be reduced or removed by the introduction of noise during training, which removes the reliance on key inputs and lessens the risk of the existence of adversarial images.

The second order coefficients dictate the effect on the function output of whether or not the values in pairs of inputs agree or disagree. For example, $\omega_5$ describes the effect of inputs 0 and 2 on the output. A positive $\omega_5$ causes the value of $\omega_5$ to be added to the function output when inputs 0 and 2 agree and to be subtracted when they disagree. Conversely, a negative value of $\omega_5$ causes the function output to increase when the values across inputs 0 and 2 differ. The second order coefficients for the digit classification network were analysed as follows. Firstly, all 300 of the second order coefficients were estimated using a sample of 50,000 random (input,prediction) pairs. Each second order coefficient relates to a pair of inputs with a connection strength. The inputs with the largest absolute summed connection strengths were identified as key pixels and the strength of their connections with each of the other 24 inputs were plotted.

Figure 9 shows an example for the output neuron associated with the digit 0. The centre pixel is the key pixel, the dark pixels are those that disagree with centre pixel when a 0 is classified, and the light pixels are those that agree with it. Note that the strongest tie of agreement is between the pixels that separately would reclassify the input as a 9 (where disagreement between these pixels is a defining feature) or a 5 (where those pixels both disagree with the centre pixel) if
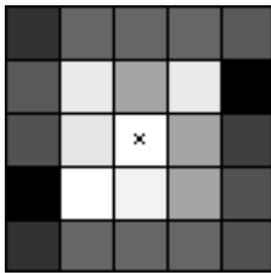
Figure 9: Second order coefficients between the centre pixel and all others. White indicates agreement, black indicates difference.



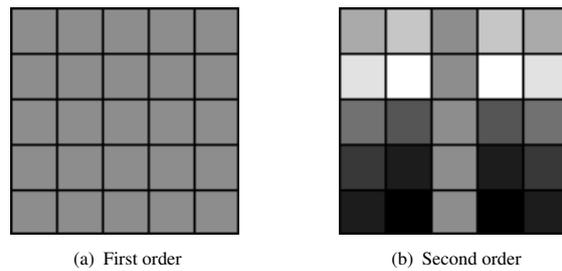(a) First order        (b) Second order

Figure 10: First and second order coefficients of a symmetry counting function. In (a), the coefficients are all zero. In (b), the shade of gray indicates a non zero second order coefficient across the two pixels with shared gray level.

their values were flipped from one to zero.

The example above was simple in the sense that the first order coefficients were sufficient to give a good insight into the function that the MLP had implemented. In the following experiment, a second order function is investigated. The function is a measure of pattern symmetry, as defined in equation 5. Figure 10(a) shows the first order coefficients of a network trained to measure the symmetry of an image. Unsurprisingly, it shows no first order coefficients of importance. Mid grey indicates values close to zero, which suggests either that the variable that corresponds to the coefficient is unimportant or that variables are involved at higher orders. The higher order coefficient values tell us which of these possibilities is true.

Figure 10(b) shows the second order coefficients of a Walsh transform of the symmetry predicting MLP. The plot is produced by finding pairs of inputs that share a non-zero second order coefficient and setting them both to the same, unique shade of gray. Note that the centre column inputs share no second order relationships and are shaded mid-gray. The others are shaded so that their gray level matches that of the inputs with which they share a non-zero second order coefficient. The depth of shade does not indicate the size of the parameter, just that a connection exists. The shading is to discriminate between input pairs.

It is clear from figure 10(b) that each input is important to the calculation of the function output, so the interpretation of the zero valued first order coefficients is that the inputs' contributions are important, but only at orders above one.

## 7 COMPARISON WITH OTHER METHODS

Recent published work in this field, such as the papers mentioned in the introduction, has concentrated on rule discovery, though what constitutes a rule is quite
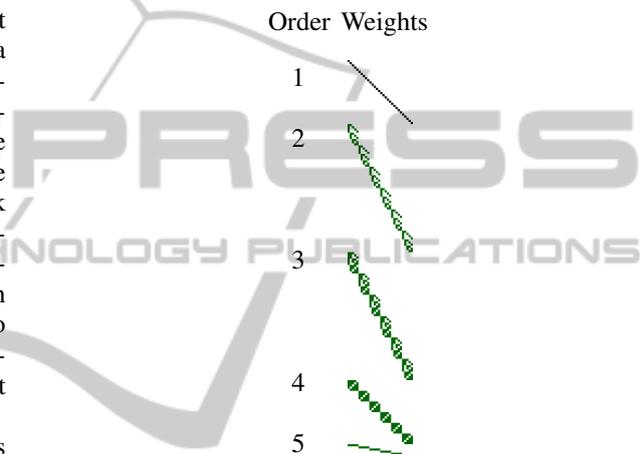


Figure 11: 5 Bit Trap Walsh Coefficients.

flexible. (Jian-guo et al., 2008), for example build a binary truth table to represent the function of the MLP. The Walsh method is a pedagogical approach, according to the definitions in (Jivani et al., 2014) as it treats the MLP as a black box. One of the advantages of the pedagogical approach is that the rules that are produced are easy to interpret.

The Walsh decomposition approach certainly aids interpretability, but it cannot be considered a rule extraction algorithm as it does not generate rules. Instead, it provides insight into the complexity of an MLP, highlighting both the level of complexity, and the variables involved. For example, in the $k$-bit trap function, it is clear from an examination of the coefficients that inputs are organised into subsets which interact within the traps, but that they are independent across traps.

One advantage of the Walsh method is that the coefficients may be easily visualised. Figure 11 shows the coefficients generated from an MLP that has learned a 5-bit trap function over 30 inputs. The figure is generated by discarding non-significant coefficients and then sorting the remaining coefficients

into combinatorial sequence so that low order coefficients are at the top of the figure. Each row of the figure represents a single coefficient as the binary equivalent of its index. For example $\omega_5$ is a second order coefficient with binary representation 101, meaning that the coefficient measures the interaction between inputs 1 and 3. Dark pixels represent connected inputs in the figure.

Another advantage of the pedagogical rule extraction approach is that it is portable across network architectures as it treats the network as a black box. The Walsh method shares this advantage. A common feature of rule extraction methods is that they accept a reduction in accuracy in return for a simpler set of rules. The rule set can be evaluated on the same test data as the MLP that generated the rules and the trade-off between accuracy and size of the rule set needs to be managed. To reproduce the functionality of the network perfectly with a rule set can require a great many rules and a large number of exceptions (or rules that apply to a very small area of input space). The Walsh method shares this limitation, but for different reasons. As the Walsh functions are a basis set, there is no function that they cannot represent, so there is no network whose behaviour cannot be perfectly reproduced. Any network with binary inputs can have its behaviour perfectly reproduced by a Walsh decomposition, but only by a full decomposition from an exhaustive sample of input,output pairs. This is possible for small networks, but infeasible for networks with large numbers of inputs. A sample of coefficients must then be calculated from a sample of data points, which will lead to an approximate representation of the MLP function.

Classification rules are generally local in that they partition a data set into subspaces that share the same output. This works well when the inputs are numeric as the conditional part of the rule can specify a range. When the inputs are discrete, as in the binary case studied here, the rules cannot partition the input space across a range. In such cases, a rule set may not be the best way to understand a function. Take the character recognition task for example, we can learn more by visualising the coefficients (even just those of low order) as shown in figure 8 than by studying a long list of rules. Walsh coefficients are global as they describe the contribution of an input or group of inputs across the entire input space. This means that it is not possible to partition the input space and so derive simple rules. Every coefficient plays a part in calculating the output from every input pattern. General statements can still be made, however, but they are of the form "When variable $x = 1$, the output increases" or "When variables $a$ and $b$ are equal, the output de-

creases". These statements can be generated directly from the coefficients.

# 8 CONCLUSIONS

An MLP trained on binary input data with either numeric or categorical output neurons can be analysed using Walsh functions. Such an analysis can reveal the relative complexity of different networks, give an insight into the way the function represented by an MLP evolves during learning and shed light on which areas of input space a network has utilised in learning that function. This understanding can help in understanding how well a network will generalise to new data and where its likely points of failure may be. An exhaustive Walsh decomposition is only possible for small networks, but a partial decomposition based on a random sample from the network's input space can still be used to gain valuable insights into the specific function learned by an MLP.

# 9 FURTHER WORK

This work has used Walsh functions as its method of complexity analysis, but other basis functions–particularly those suitable for real valued inputs–are also worthy of investigation. As the analysis is not designed to reconstruct the function, merely to shed light on its structure in a human readable form, it should be possible to use an information theoretic measure of interaction such as mutual entropy.

The method provides a useful measure of network complexity that is not based on the number of weights in the network. Training methods that favour simple models over more complex ones often use parameter counts (in the case of MLP, the weights) as a measure of complexity. For example, minimum description length (MDL) methods are often based on parameter counts, but might usefully be adapted to account for other types of complexity such as that described here. The Walsh analysis reveals that two networks of equal size do not necessarily share an equal complexity. The relationship between network complexity and network size is an interesting field of study in its own right. Of course, this analysis is not restricted to use with MLPs. Any regression function may be used, but it is well applied to MLPs as they are difficult to analyse in terms of the structure of their weights alone.

The number of Walsh coefficients to consider grows exponentially with the number of inputs to the network, so it is not possible to exhaustively calculate every possible one in a large network. For networks

that contain key interactions at a number of different higher orders, the task of finding the significant coefficients becomes a great problem. Work on heuristics for finding the significant high order coefficients in a sparse coefficient space is ongoing. One approach is to build a probabilistic model of the importance of different neurons and connection orders and sample coefficients from that model. As more coefficients are found, the quality of the model improves and allows the faster discovery of others.

## REFERENCES

Augasta, M. and Kathirvalavakumar, T. (2012). Rule extraction from neural networks - a comparative study. pages 404–408. cited By (since 1996)0.

Bartlett, E. B. (1994). Dynamic node architecture learning: An information theoretic approach. *Neural Networks*, 7(1):129–140.

Baum, E. B. and Haussler, D. (1989). What size net gives valid generalization? *Neural Comput.*, 1(1):151–160.

Beauchamp, K. (1984). *Applications of Walsh and Related Functions*. Academic Press, London.

Castillo, P. A., Carpio, J., Merelo, J., Prieto, A., Rivas, V., and Romero, G. (2000). Evolving multilayer perceptrons.

Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14(2):179–211.

Gorman, R. P. and Sejnowski, T. J. (1988). Analysis of hidden units in a layered network trained to classify sonar targets. *Neural Networks*, 1(1):75–89.

Hruschka, E. R. and Ebecken, N. F. (2006). Extracting rules from multilayer perceptrons in classification problems: A clustering-based approach. *Neurocomputing*, 70(13):384 – 397. Neural Networks Selected Papers from the 7th Brazilian Symposium on Neural Networks (SBRN '04) 7th Brazilian Symposium on Neural Networks.

Jian-guo, W., Jian-hong, Y., Wen-xing, Z., and Jin-wu, X. (2008). Rule extraction from artificial neural network with optimized activation functions. In *Intelligent System and Knowledge Engineering, 2008. ISKE 2008. 3rd International Conference on*, volume 1, pages 873–879. IEEE.

Jivani, K., Ambasana, J., and Kanani, S. (2014). A survey on rule extraction approaches based techniques for data classification using neural network. *International Journal of Futuristic Trends in Engineering and Technology*, 1(1).

Kamimura, R. (1993). Principal hidden unit analysis with minimum entropy method. In Gielen, S. and Kappen, B., editors, *ICANN 1993*, pages 760–763. Springer London.

Krogh, A. and Vedelsby, J. (1994). Neural network ensembles, cross validation, and active learning. In *NIPS*, pages 231–238.

Kulluk, S., Özbakir, L., and Baykasoğlu, A. (2013). Fuzzy difaconn-miner: A novel approach for fuzzy rule extraction from neural networks. *Expert Systems with Applications*, 40(3):938 – 946. FUZZYSS11: 2nd International Fuzzy Systems Symposium 17-18 November 2011, Ankara, Turkey.

Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1986). Parallel distributed processing: Explorations in the microstructure of cognition, vol. 1. chapter Learning Internal Representations by Error Propagation, pages 318–362. MIT Press, Cambridge, MA, USA.

Saad, E. and Wunsch II, D. (2007). Neural network explanation using inversion. *Neural Networks*, 20(1):78–93. cited By (since 1996)22.

Sanger, D. (1989). Contribution analysis: A technique for assigning responsibilities to hidden units in connectionist networks. *Connection Science*, 1(2):115–138.

Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I. J., and Fergus, R. (2013). Intriguing properties of neural networks. *CoRR*, abs/1312.6199.

Uphadyaya, B. and Eryurek, E. (1992). Application of neural networks for sensor validation and plant monitoring. *Neural Technology*, (97):170–176.

Walsh, J. (1923). A closed set of normal orthogonal functions. *Amer. J. Math*, 45:5–24.

Weigend, A. S., Huberman, B. A., and Rumelhart, D. E. (1992). Predicting Sunspots and Exchange Rates with Connectionist Networks. In Casdagli, M. and Eubank, S., editors, *Nonlinear modeling and forecasting*, pages 395–432. Addison-Wesley.

Widrow, B. and Lehr, M. (1990). 30 years of adaptive neural networks: perceptron, madaline, and backpropagation. *Proceedings of the IEEE*, 78(9):1415–1442.

Yao, X. (1999). Evolving artificial neural networks. In *Proceedings of the IEEE*, volume 87, pages 1423–1447.