ESANN'1996 proceedings - European Symposium on Artificial Neural Networks
Bruges (Belgium), 24-25-26 April 1996, D-Facto public., ISBN 2-9600049-6-5, pp. 161-166

Investigating Lexical Access Using Neural Nets

Malti Patel

Macquarie University, Department of Computing,
North Ryde, Sydney, NSW 2109, Australia

Abstract.
Psychological theories of reading have been around for many years. With the advent of neural nets, we can at last start to test these theories since they are based on the workings of the brain. In this paper, I will show how a neural net has been used to test theories regarding how the meanings of a word are accessed. The results give a new angle to this area in that they suggest that meanings may be accessed in a way not expressed by any of the theories. This shows how neural nets can be used at a practical level to test psychological theories and how they can sometimes bring to light new possibilities.

1. Introduction

Many experiments have been performed to determine how people store the spelling, pronunciations and meanings of words, and how these representations are used by the processes involved in reading, determining whether something is a word, and so on. Data from these experiments have previously been used to hypothesise exactly what is happening in the brain. There is no concrete evidence to support these hypotheses. With the advent of neural networks, we can at last start to test these theories.

In this paper, we will be using a neural net to examine how the meanings (senses) of a word are accessed when a word is read. For example, when you see the word BALL, do you just think of something that is round and bounces, or do you also, at some level, think of the meaning associated with dancing, i.e. "going to a BALL".

2. Theories of Lexical Access

Lexical access simply means accessing words stored in the brain. There are three main theories as to how the meaning(s) of a word is accessed. We will look at these briefly. There is evidence to support all theories, see Garnham [2] for a review.

I would like to acknowledge the help received from Max Coltheart in Behavioural Sciences at Macquarie University which made this work possible.

ESANN'1996 proceedings - European Symposium on Artificial Neural Networks
Bruges (Belgium), 24-25-26 April 1996, D-Facto public., ISBN 2-9600049-6-5, pp. 161-166

- **Exhaustive Access Theory**

  This states that all the meanings of a word are accessed when we see a word. A choice is then made according to the context.

- **Ordered Access Theory**

  This states that meanings are accessed serially in order of their frequency. Each meaning is matched against the context until the correct meaning is found. This means that when the most frequent meaning of a word is encountered, the word is effectively acting like an unambiguous word.

- **Context-Guided Access**

  This states that only the meaning associated with the relevant context is accessed.

From the outside, the first theory seems implausible since most people would say that other meanings of the same word were not noticed. For example, seeing the word BALL, most people would say that they never thought of the meaning associated with dancing. Theory 3 seems the most plausible since most people would think that they only associate one meaning whenever a word is encountered.

3. The Dual-Route Reading Model (DRM)

The DRM [1] is a theory of reading aloud which has been implemented using a localist neural network. It currently pronounces about 8000 words. Figure 1 shows the structure of this model with the various excitatory and inhibitory links. The spelling of a word is input in terms of its letter features. The model has to pronounce the word using the two routes shown.

The grapheme-phoneme rule route relies on rules to pronounce a word. For example, one rule would say that when the letters "oo" are encountered, they should be pronouned as in "mood" not as in "flood".

The other route has a specific entry for each word recognized by the model. For example, the spelling of BALL would be represented by one unit in the *visual word-detectors* component, its pronounciation would be represented by one unit in the *phonological unit set.* One unit is used to represent each sense (meaning) of BALL in the "senses" component. The features used to define each sense (e.g. round, bounces) are stored in the *semantic features* component. There are connections between the sense of a word and the semantic features used to define it. The semantic features were obtained from a psycholinguistic on-line dictionary called WordNet, see [3] for more details. This second route therefore uses the spelling and meaning(s) of a word to pronounce it.

When adding the semantic component, the aim was that the system output the sense with the highest frequency of the input word. For the highest frequency sense to win, there must obviously be information about the frequency values of the various senses of a word. Such data is fairly difficult to obtain

ESANN'1996 proceedings - European Symposium on Artificial Neural Networks
Bruges (Belgium), 24-25-26 April 1996, D-Facto public., ISBN 2-9600049-6-5, pp. 161-166

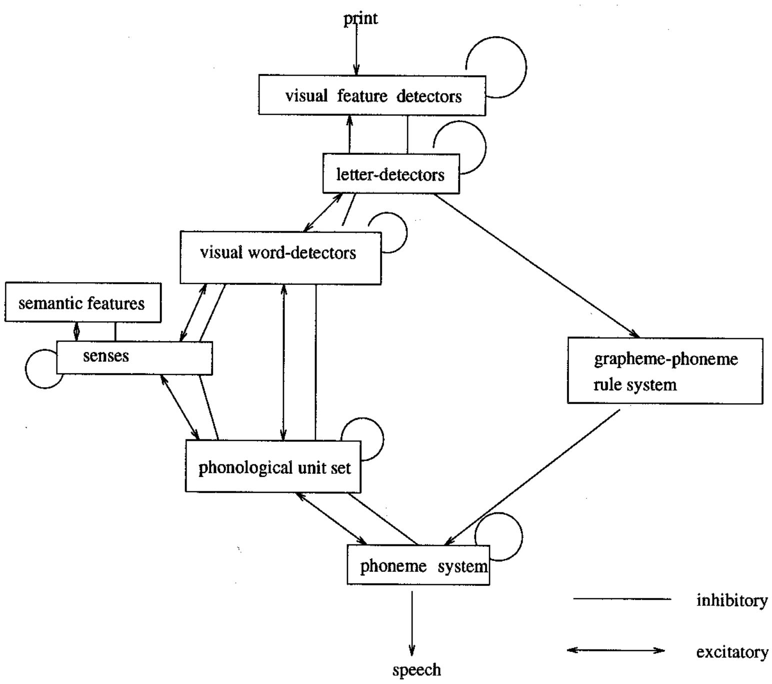

Figure 1: The Dual-Route Reading Model

since frequency values are usually given for a whole word and not its separate meanings. However, we were able to obtain data for the senses of 450 words from (Twilley, Dixon et al) [4]. Thus, currently we have the system running on 450 words when incorporating the semantic component.

The sense with the highest activation wins. The frequency factor determines how much activation a sense receives, i.e. a high frequency value leads to more activation being received by a sense. Since we want a winner take all situation, the senses of a word compete against each other by means of lateral inhibition. Hence, for a word, the sense with the most activation can send more inhibition to the other senses of that word.

4. Results from the Model

In this section we will illustrate the activation of the sense units for different inputs.

Figure 2 shows the activations of the two senses of SIGN when that word is input. The high frequency sense wins with no problem. The low frequency sense is also activated but its activation starts to decrease eventually. This was found with most senses - all senses were activated to some degree, but whereas the activation of the high frequency (HF) sense kept increasing, that of the low frequency sense (LF) started to decrease eventually.

ESANN'1996 proceedings - European Symposium on Artificial Neural Networks
Bruges (Belgium), 24-25-26 April 1996, D-Facto public., ISBN 2-9600049-6-5, pp. 161-166

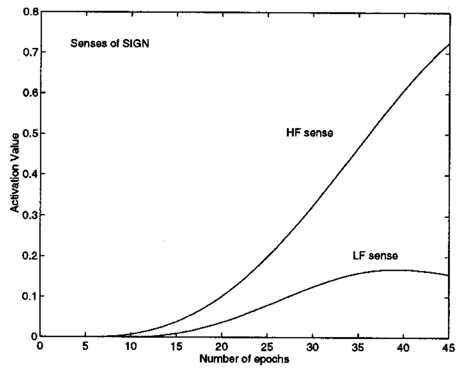

Figure 2: Activation of the Senses of SIGN

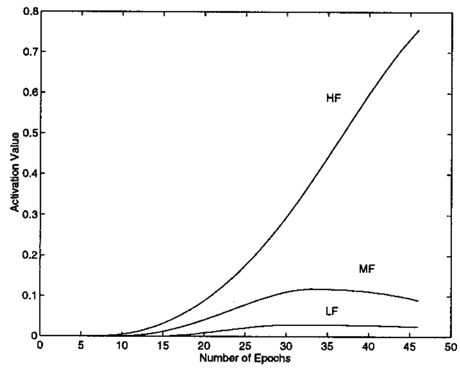

Figure 3:Activation of the Senses of DRAFT

The lateral sense inhibition is an important factor here. As the activation of the HF sense increases, this starts to dampen the activation of any LF senses because they receive a large amount of inhibition. This also means that the HF sense receives hardly any inhibition from the LF senses since their activation value will remain low.

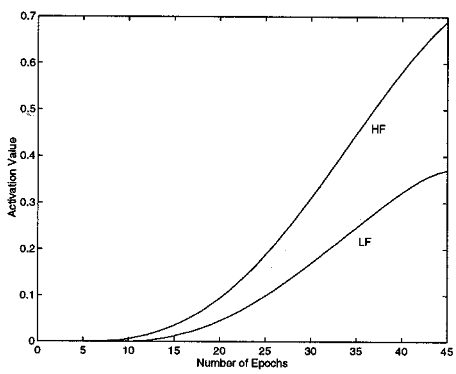

Figure 4: Activation of the Senses of GRAIN

Figure 3 shows the activations of the three senses of DRAFT. As seen, all three are activated, however slight that may be. The medium frequency (MF) and low frequency senses decrease eventually.

Figure 4 shows the activations of the two senses of GRAIN. In this case, the LF sense is still increasing when the word is pronounced, even though its activation is much lower than that of the HF sense.

Overall, the effect was that either the LF sense decreased eventually or kept increasing. In some cases, the LF senses received no activation at all.

Analysing the frequency values of the senses gives an insight into the results above. In some cases, the frequency values of the senses of a word may be far apart. For example, the HF sense may have a value of 0.8 and the LF sense may have a value of 0.1 (out of say 1.0). These are known as *polarized* senses. Intuitively, it seems highly likely that in this case, it is very difficult for the LF sense to get any activation and even if it does, to have that activation increase

ESANN'1996 proceedings - European Symposium on Artificial Neural Networks
Bruges (Belgium), 24-25-26 April 1996, D-Facto public., ISBN 2-9600049-6-5, pp. 161-166

all the time. This accounts for senses such as SIGN.

Alternatively, the frequency values of some senses may be close together, e.g. the HF sense may have a frequency value of 0.7 and the LF sense may have a value of 0.5. These are *equiprobable* senses and it seems likely that in this case the activations of both senses will increase continually. In the DRM, the activations of both senses will be similar and thus the lateral inhibition will not have a large effect. This accounts for words such as GRAIN where the activation of the LF sense never decreases.

Overall, it seems as if the *difference* in the frequency values of individual senses have a great effect as to how much activation each sense receives.

5. Relating Results to the Theories of Lexical Access

Comparing the theories of lexical access against the results we can offer new insights as to what may be occurring. It may be that the different theories arise because of the types of senses, i.e. polarized or equiprobable, as opposed to there actually being different ways of accessing meanings.

The exhaustive access theory states that all meanings are accessed. As seen, in the case of equiprobable senses, this may well be the case. The activation of the meanings are all highly active and therefore it could be said that all meanings are accessed.

The ordered access theory states that meanings are accessed in order of frequency. This could well be true of polarized senses where the frequency of the LF sense is so far from the HF meaning that it could well be viewed that only the HF sense has been accessed. The activation of the LF sense is so low that it appears as if it has not been accessed.

Contextually-guided access is plausible when reading a sentence or text if we view words are being connected according to their semantics. This means, that seeing a word like NURSE would automatically send activation to semantically associated words such as DOCTOR. In our model, we have implemented this via the WordNet definitions of words which are stored in the semantic features component. For example, in WordNet, the related meanings of HEAD and HAND share some of the same semantic features. Thus, we have an associative semantic representation such that if HEAD is input, it will also activate the related sense of HAND. This means, that if related words are input, the activation of the semantic features will lead to semantically related meanings being activated.

As far as accessing meanings is concerned, it may be that people are "conscious" of all meanings which are equiprobable and of the HF meanings of polarized meanings. In the latter case, the LF meaning may never be active enough be deemed as being noticed. This would explain how the exhaustive and ordered access theories arise.

When reading a sentence or some text, it is easy to imagine that activation is being sent to semantically related words. For example, say that word X is semantically related to the LF sense of word Y. Say word X appears on the

ESANN'1996 proceedings - European Symposium on Artificial Neural Networks
Bruges (Belgium), 24-25-26 April 1996, D-Facto public., ISBN 2-9600049-6-5, pp. 161-166

screen. Activation would automatically be sent to the LF sense of word Y. Now when word Y appears on the screen, there is a greater probability that its LF sense would be more active than the HF sense. Hence, we get evidence for contexually-guided access since in this case it is not the frequency of a sense which matters but how much activation semantically related senses have received.

Overall, it seems that it is the frequency of a sense which determines how it is accessed and not that all senses are accessed in the same manner as stated by the different access theories. As shown above, all three theories are plausible if we examine the frequencies and activation values of the senses of a word.

6. Conclusion

We have shown how a neural net implementation of a psychological theory of reading can be used to examine different theories of lexical access. The neural net is modelled on a theory which has much well-established data to supprt it, therefore, it is a plausible model of what occurs in the brain when people read.

The neural net has shown at a very low level how the meanings of a word may be accessed when we read. This is not possible when devising experiments to test how people access meanings. The former lead to theories which our results have shown may be based on the wrong assumption. Namely, people do not access meanings in a certain way, but that the frequency of a meaning determines how it will be accessed. The frequency indicates how much activation that meaning will receive and hence if it will be "noticed" by the person reading.

Neural nets can be used in cognitive modelling to test psychological theories at a practical level. As such, they have a large role to play since there are many theories which can be tested. As seen, they can lead to interesting new insights into how language is stored and accessed by the brain.

References

[1] M. Coltheart, B. Curtis, P. Atkins, and M. Haller. Models of reading aloud: Dual-route and parallel-distributed-processing approaches. *Psychological Review*, 100:589–608, 1993.

[2] Alan Garnham. *Psycholingustics: Central Topics.* Methuen and Co, 1985.

[3] G.A. Miller, R. Beckwith, R. Fellbaume abd C. Gross, D. Miller, and R. Tengi. Five papers on wordnet. Csl report, Princeton University, 1990.

[4] L.C. Twilley, D. Dixon, D. Taylor, and K. Clark. University of norms of relative meaning frequency for 566 homographs. *Memory and Cognition*, 22:111–126, 1994.