FNLP Tutorial 5

1 Lexical semantics

1. Consider the words "pie" (dish baked in pastry-lined pan often with a pastry top), "relationship" (human relationship), "universe" (everything that exists anywhere) and "garden" (ground where plants are cultivated). At first glance, these words seem to have very little in common. Use Wordnet (http://wordnetweb.princeton.edu/perl/webwn) to follow the chain of hypernyms for these words, to figure out which one is the odd one out. Report the chain of hypernyms and use them to motivate your answer. You follow the chain of hypernyms by clicking 'S' next to a word sense, and looking at its 'direct hypernym'.

2. Describe the relationship between the following word pairs. Where there are multiple plausible word senses for the relation, we disambiguate this. How far do you get with the main relationships in WordNet (hyponym/hypernym, meronym, antonym)? Characterise the relationships that are not covered by WordNet and give them a name; try to invent relationships that could potentially be used for many other word pairs.

   - bike - vehicle
   - bike - sports
   - book - novel
   - book - library (a room where books are kept)
   - fish (the flesh of fish used as food) - chips (a thin crisp slice of potato fried in deep fat)
   - knife - drawer
   - knife - cut

2 Word vectors

Consider the following two-dimensional word vectors that encode animals. They are projections of 300-dimensional word vectors, that were trained using data from all over the web.

$$\overrightarrow{\text{dog}} = \begin{bmatrix} 1.9 \\ 2.07 \end{bmatrix}, \overrightarrow{\text{puppy}} = \begin{bmatrix} 2.04 \\ 1.82 \end{bmatrix}, \overrightarrow{\text{puppies}} = \begin{bmatrix} 2.15 \\ 1.77 \end{bmatrix}, \overrightarrow{\text{cat}} = \begin{bmatrix} 1.63 \\ 1.87 \end{bmatrix}$$
$$\overrightarrow{\text{kitten}} = \begin{bmatrix} 1.68 \\ 1.76 \end{bmatrix}, \overrightarrow{\text{snake}} = \begin{bmatrix} 0.25 \\ 1.44 \end{bmatrix}, \overrightarrow{\text{monkey}} = \begin{bmatrix} 0.02 \\ 1.56 \end{bmatrix}, \overrightarrow{\text{monkeys}} = \begin{bmatrix} -0.04 \\ 1.59 \end{bmatrix}$$

1. Visually inspect the word vectors. Distances in the vector space capture word similarities (albeit strongly simplified when using only two dimensions). Why would 'snake' (a reptile) be closer to 'monkey' (a mammal) compared to the remaining animals (that are also mammals)? Please refer to the distributional hypothesis and the way word vectors are constructed in your answer.

2. Word vectors can capture relationships between words, as discussed in J&M Section 6.10, 3rd edition. An example of such a relationship is an analogy, such as Edinburgh is to Scotland as Canberra is to . . . ? When we have word vectors, we can use the parallelogram method to

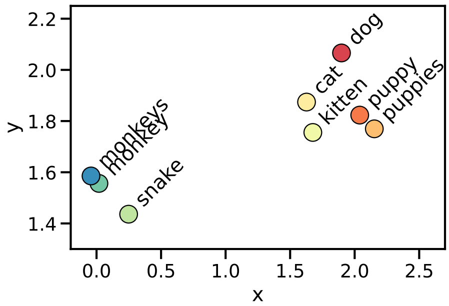

solve analogies: by subtracting $\overrightarrow{\text{edinburgh}}$ from $\overrightarrow{\text{scotland}}$ and adding $\overrightarrow{\text{canberra}}$. You would pick the word for which the word vector has the smallest distance to the resulting vector. J&M represent the procedure as follows for the analogy a:b::a*:b* (a is to b as a* is to b*):

$$\vec{\text{b}}* = \text{argmin distance}_{\vec{\text{x}}}(\vec{\text{x}}, \vec{\text{b}} - \vec{\text{a}} + \vec{\text{a}}*) \tag{1}$$

Using Euclidean distances, compute $\vec{\text{b}}*$ for:

- $\vec{\text{a}} = \overrightarrow{\text{dog}}$, $\vec{\text{b}} = \overrightarrow{\text{puppy}}$ and $\vec{\text{a}}* = \overrightarrow{\text{cat}}$
- $\vec{\text{a}} = \overrightarrow{\text{puppy}}$, $\vec{\text{b}} = \overrightarrow{\text{puppies}}$ and $\vec{\text{a}}* = \overrightarrow{\text{monkey}}$

3. Retrieve the word that is the most similar to ‘monkey’ using cosine similarity, euclidean distance and the dot product. Where do the different metrics disagree, and how does this relate to their definition?

4. Explain how Euclidean distance, the dot product and the cosine similarity are related.

3 Feed-forward neural networks

Consider a two-layer neural network with the topology visualised below, with the corresponding weights and bias values in the table. The hidden layer is followed by a non-linear function: the ReLU. The output layer is followed by a non-linear function too: the softmax. Read up on those functions and how to work with feedforward neural networks in sections 7.1 to 7.3 from J&M (only available in the 3rd edition!).

The network can be used for simple classification for three output classes. An input (consisting of features x_1 and x_2) belongs to one of the three classes, and you will classify an example input.

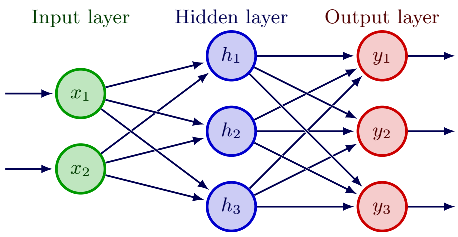

$w_{x_1,h_1} = 0.23$	$b_{h_1} = 0.13$	$w_{h_1,y_1} = 0.02$	$b_{y_1} = 0.31$
$w_{x_1,h_2} = -2.90$	$b_{h_2} = 0.90$	$w_{h_1,y_2} = -0.13$	$b_{y_2} = -0.18$
$w_{x_1,h_3} = -0.45$	$b_{h_3} = 0.24$	$w_{h_1,y_3} = 0.47$	$b_{y_3} = -0.56$
$w_{x_2,h_1} = 1.10$		$w_{h_2,y_1} = -0.88$	
$w_{x_2,h_2} = 0.05$		$w_{h_2,y_2} = 1.74$	
$w_{x_2,h_3} = -0.01$		$w_{h_2,y_3} = -0.09$	
		$w_{h_3,y_1} = -0.55$	
		$w_{h_3,y_2} = 0.72$	
		$w_{h_3,y_3} = -2.22$	

1. Compute the class an input with $x_1 = 1.50$, $x_2 = 3.11$ would belong to. Show the intermediate computations, not just the final class.

2. Now imagine that you want to perform classification, but one input can belong to multiple classes. For example, when classifying a sentence with an emotion, that sentence can capture both anger and despair. To enable multi-class classification in this network, what adaptation would you make to its structure or the non-linear functions it uses?