2010 Paper 9 Question 9

Natural Language Processing

The following is a fragment of an annotated corpus that could be used for training and testing a pronoun resolution algorithm.

```
<s>Adapting <e id="1" pron="f">this first verse</e>,
<e id="2" pron=f>Burns</e> did <e id="3" pron="f">something
unforeseen</e>.<s><e id="4" pron="t">He</e> obscured and
mystified <e id="5" pron="f"> <e id="6" pron="t">its</e>
sentiments</e>

Key: 4=2, 6=1
```

`<e></e>` delimits a referring expression, `pron` indicates whether the expression is a pronoun to be resolved, and `<s>` marks sentence boundaries. The key gives the antecedents of the pronouns.

(a) Using this fragment as an example, illustrate how pronoun resolution can be treated as supervised classification using a Naïve Bayes classifier. You should include details of possible features, explaining why they are relevant and outlining how they might be derived where appropriate. Details of the Naïve Bayes classifier are not required.   [15 marks]

(b) Why might better performance be achieved in pronoun resolution if a discourse model that records how other pronouns have been resolved is utilised?
                                                                      [3 marks]

(c) What is meant by the term *baseline* in the context of NLP evaluation? Suggest a possible baseline algorithm for the pronoun resolution task.   [2 marks]