School of Computing and Information Systems
The University of Melbourne
COMP90042 WEB SEARCH AND TEXT ANALYSIS (Semester 1, 2019)

Workshop exercises: Week 8

Discussion

1. What is **Information Extraction**? What might the "extracted" information look like?
   (a) What is **Named Entity Recognition** and why is it difficult? What might make it more difficult for persons rather than places, and *vice versa*?
   (b) What is the **IOB** trick, in a sequence labelling context? Why is it important?
   (c) What is **Relation Extraction**? How is it similar to NER, and how is it different?
   (d) Why are hand–written patterns generally inadequate for IE, and what other approaches can we take?
2. What is **Question Answering**, and how is it related to **Information Retrieval** and Information Extraction?
   (a) What is **semantic parsing**, and why might it be desirable for QA? Why might approaches like NER be more desirable?
   (b) What might be the main steps for answering a question for a QA system?

Programming

1. NLTK comes with a pre-trained named entity **chunker** `ne_chunk`, which takes a tagged sentence as input, and outputs a tree:

```
>>> print(nltk.ne_chunk(nltk.corpus.treebank.tagged_sents()[11]))
(S
  Dr./NNP
  (PERSON Talcott/NNP)
  led/VBD
  a/DT
  team/NN
  of/IN
  researchers/NNS
  from/IN
  the/DT
  (ORGANIZATION National/NNP Cancer/NNP Institute/NNP)
  ...
```

   Read up on how to traverse an `nltk.tree.Tree` object, and then convert the tree into a (flat) IOB-representation.

Catch-up

- What is **POS tagging**, and what are some common methods for applying it?
- What is a **Named Entity**?
- What kinds of **relations** can exists between tokens ("words")? Constituents? Sentences? Documents?
- What is **parsing** and how is it different to **tagging**?

Get ahead

- Try using the `nltk.sem.extract_rels()` to extract a set of relations from the collection `nltk.corpus.ieer.parsed_docs()`, and then write a system that can answer simple questions like: "Where is [the] Bastille Opera?"