

NLP & IR

Examination

Dell Zhang

Birkbeck, University of London

2020/21

Exam Paper

- **$6 \times 15 + 10 = 100$**
 - There are seven questions in total
 - 6 with 15 marks each.
 - 1 with 10 marks
 - Students should answer **all** of them.
 - The exam paper is not prior-disclosed.

Exam Paper

- ~~• Bookwork Questions: Factual Knowledge
($\approx 10\%$)
 - ~~– They should be able to be answered in a few sentences.~~
 - ~~– If you give all the relevant factual points, you'll get all the marks. Don't overdo it.~~~~
- Technical Questions: Problem Solving
($\approx 90\%$)

Materials

- Lecture Slides
- Textbook Chapters
 - All topics are covered by the Stanford IR textbook and the Stanford NLP textbook.
- Classwork and Coursework
- Past Exam Papers
 - Some are out-of-date and out-of-scope now
 - Similar styles of questions this year

What's New

- This module used to be called Information Retrieval and Organisation (IRO)
- Compared to the previous years
 - n-gram language models, perplexity, ...
 - PPMI, 1st-order/2nd-order co-occurrences, ...
 - word2vec and neural language models

Tips

- Manage Your Time
- Software Tools
 - Adobe Acrobat Reader
 - Edit / Take a Snapshot
 - Microsoft Word
 - Plain Text Editor
 - e.g., Notepad++, Sublime Text 3, ...
 - Online Calculator
- Copy and Paste
 - Exploit the Patterns
 - Notice the Differences
- Module Materials at Your Fingertips
 - Search
 - Scripts