




Artificial Intelligence Programming

Information Retrieval

Chris Brooks

Department of Computer Science
University of San Francisco



Processing text

- Now that we know a little bit about how to consider and compare different states, we'll think about a problem with a harder representation.
- English-language text
 - Specifically, webpages
- We'll look at several different approaches to automatically summarizing and understanding documents.

Difficulties

- What makes working with natural language hard?