Natural Language and AI

Prof. Neil C. Rowe

Computer Science Department
Naval Postgraduate School

http://faculty.nps.edu/ncrowe

Fall 2019

Natural language processing in AI

- A distinctive feature of human intelligence is communication by languages.
- Language processing has been an important goal for AI since it is essential to human intelligence.
- Humans communicate much more easily by language than by online forms.
- Progress has been slow but steady for 60 years.
- Natural-language processing (NLP) is mainly useful for unstructured text.
- Applications: (1) communicating with automated systems, (2) summarizing text, (3) extracting clues from text, (4) translating text.

Subareas of natural-language processing

- Speech understanding: Signal processing and segmentation into sounds (phonemes).
- Word recognition: Using a dictionary.
- Morphology: Understanding punctuation and word suffixes and prefixes.
- Grammar (syntax): Recognizing word sequence structure.
- Semantics: Assigning meanings to structures.
- Discourse analysis: Assigning meanings to larger word structures like paragraphs.
- Translation: Converting from one language to another.

Automated speech understanding today

- Automated speech understanding is increasingly commercial in phone-call management, automated dictation transcription, and now in digital interfaces like Alexa, Siri, and Google Assistant.
- Most methods used today were developed 40 years ago – but only became practical in the last 10 years as digital devices became fast enough.
- Most systems match a large database of possible pronunciations of words by different speakers ("shallow" processing). Better accuracy is possible by knowing common sequences of words, or by training for a particular speaker.

Example speech pattern: "phonetician"

Speech signals are quite noisy, even without background noise.

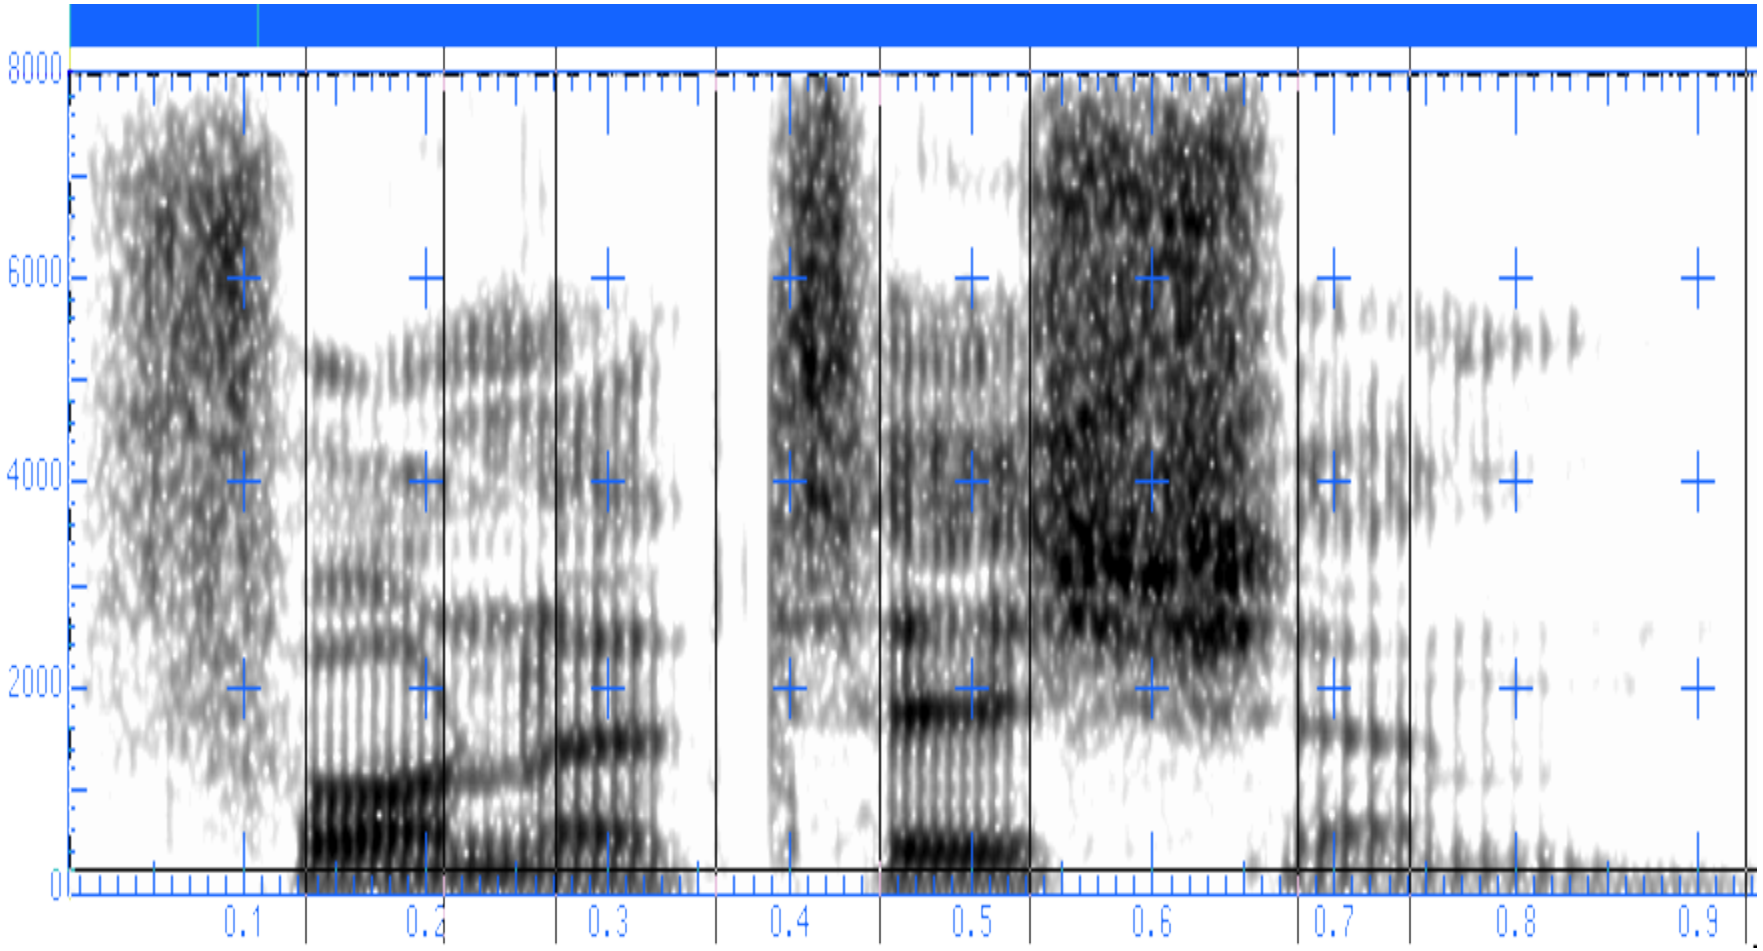

Frequency in cycles per second

Seconds

"Shallow" (statistical) strategy for natural-language processing, both speech and text

Look for words, pairs of successive words, triples of successive words, etc. Most AI approaches and Google do this. Use probabilistic methods to:

- Recognize words from a small set of possibilities, e.g. spoken numbers.
- Find text matching some keywords.
- Classify some text (e.g. recognize spam, authors, emotions, or deception).
- Learn to associate co-located words.
- Translate text based on previous translations of parts of the text.

Some words associated with deception

- Decreased use of “I”, “we”, “my”
- Decreased use of “except”, “unless”, “without”, “however”
- Increased use of “hate”, “dislike”, “ignore”, “lose”
- Increased use of “move”, “go”, “carry”, “take”

Count these in some text to estimate its degree of deception.

Example shallow natural-language analysis

- "Chinese powerboat reported at 30N 170E at 1300 heading east at 20 knots leaking oil."
- Like Google, we can recognize and index "report", "powerboat", "heading", "knots", and "oil".
- We can also guess "30N 170E" is a position, "1300" as a time, and "carrier" as the object having that data. We can enter this data into a ship database.
- We can also learn useful associations between "Chinese" and "powerboat", and between "powerboat" and "leaking".

Example Spanish translations from Google Translate

These were done in less than a second.

- Network security -> seguridad de red
- Software suite -> paquete de programas
- Network-security software suite -> paquete de software de seguridad de red
- Purchase our new software -> Compre nuestro nuevo software
- Purchase our new network-security software suite -> Compre nuestro nuevo paquete de software de seguridad de red
- Purchase our new network-security software suite in the next month -> Compre nuestro nuevo paquete de software de seguridad de red en el próximo mes
- Purchase our new network-security software suite today -> Compre hoy nuestro nuevo paquete de software de seguridad de red

Deep ("linguistic") strategy for natural-language understanding

(Not to be confused with "deep" learning.) Tries to understand the full meaning of text:

- Requires a full parse and full semantic interpretation.
- Requires analysis of goals and intentions.
- Based on the field of linguistics.
- Deep understanding is needed when the exact content and context matters, as in legal reasoning.
- Deep methods will eventually surpass shallow methods since language changes only slowly.

Example of deeper structure analysis

- "Powerboat reported at 30N 170E at 1300 heading east at 20 knots."

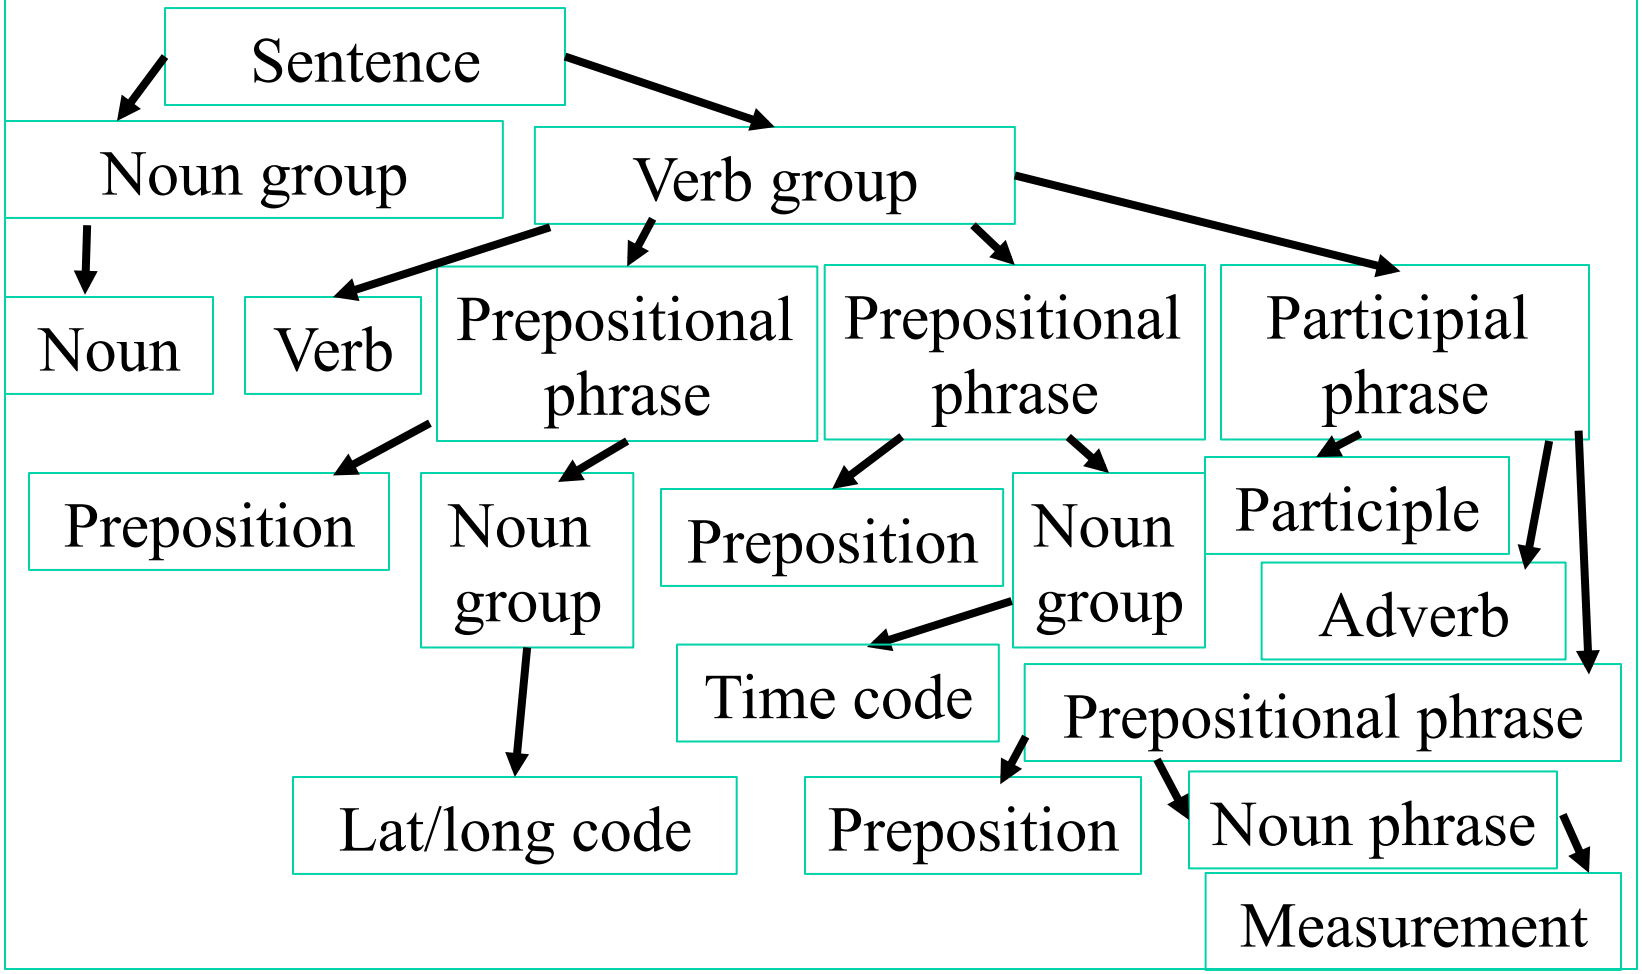

Example of deeper semantic analysis

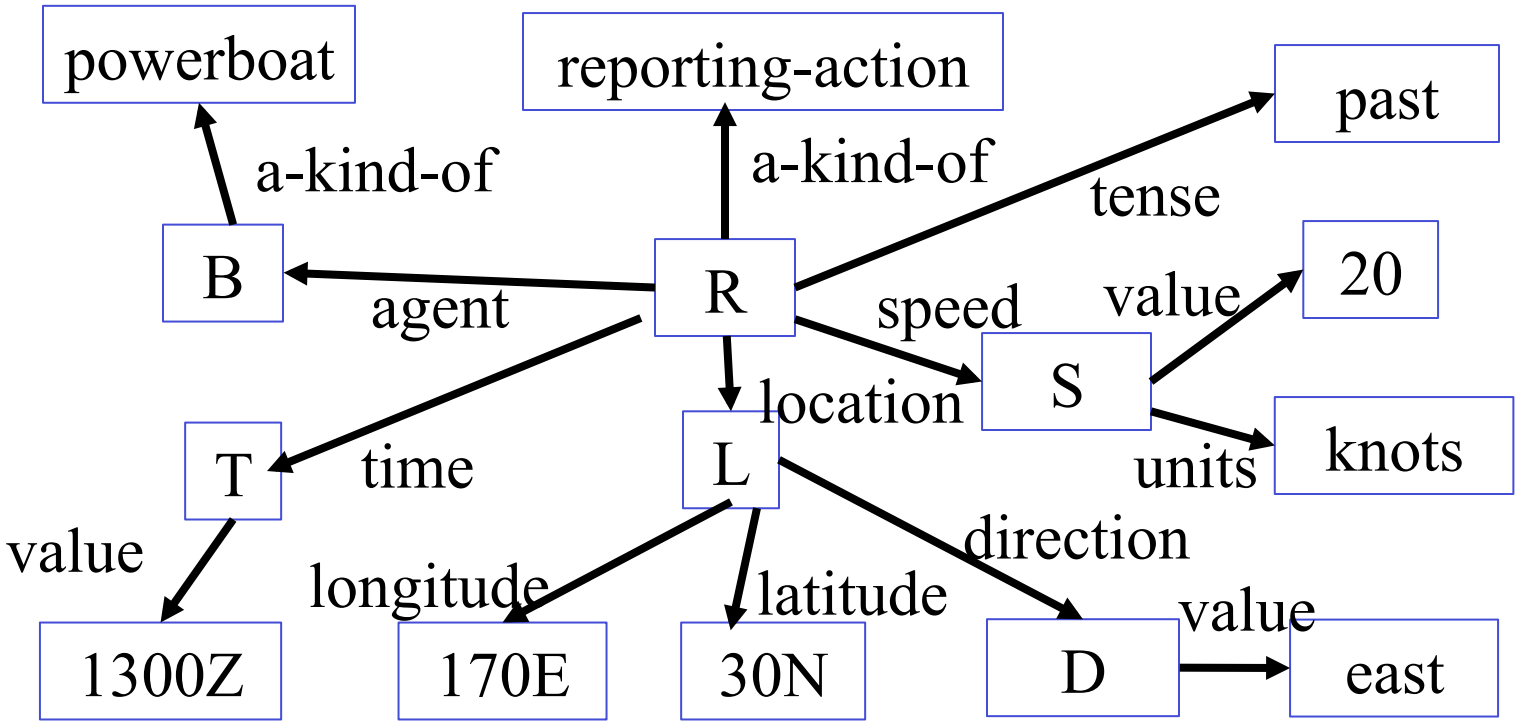

"Powerboat reported at 30N 170E at 1300 heading east at 20 knots.". R is a reporting action, B is a boat, S is a speed, L is a location, T is a time.

Example of deeper discourse analysis

- "Chinese powerboat reported at 30N 170E at 1300 heading east at 20 knots leaking oil. Nearest port is Base 83, but it is not a full maintenance facility."
- Analysis needs to infer from world knowledge:
  - "It" means port.
  - The powerboat is in trouble, from knowledge of what "leaking oil" means.
  - It needs maintenance.
  - It must get maintenance at a port.
  - Leaking oil depletes ship functionality and needs to be fixed soon and nearby.

Key problems in natural-language processing (1)

- Speech is a noisy signal which requires some guesses and use of context to decipher.
- The same word can serve multiple grammatical functions, and you can't understand a sentence unless you figure each word's purpose. Solution: Part-of-speech taggers.
- Words have multiple meanings: An average dictionary entry gives 4 meanings for a word.

Key problems in natural-language processing (2)

- Words relate to one another in many ways. In English, prepositions and successive-noun pairs require complex reasoning to determine relationships of words.
- "Anaphoric references" like pronouns and vague nouns require context analysis to understand.
- Understanding also requires theories of causation, how arguments are constructed, and social interactions between people ("speech acts").
- People expect high accuracy in natural-language conversations: 95% accuracy is not enough.