Natural Language and AI

Prof. Neil C. Rowe

Computer Science Department
Naval Postgraduate School

http://faculty.nps.edu/ncrowe

Summer 2021

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.
- 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 in time-critical situations like flying aircraft, (2) categorizing documents like reports, (3) extracting key concepts from documents like intelligence, (4) translating text or speech.

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 the usual sequences of words, or by training for a particular speaker.

Example speech pattern: "phonetician"

Speech signals are quite noisy, even without background noise.

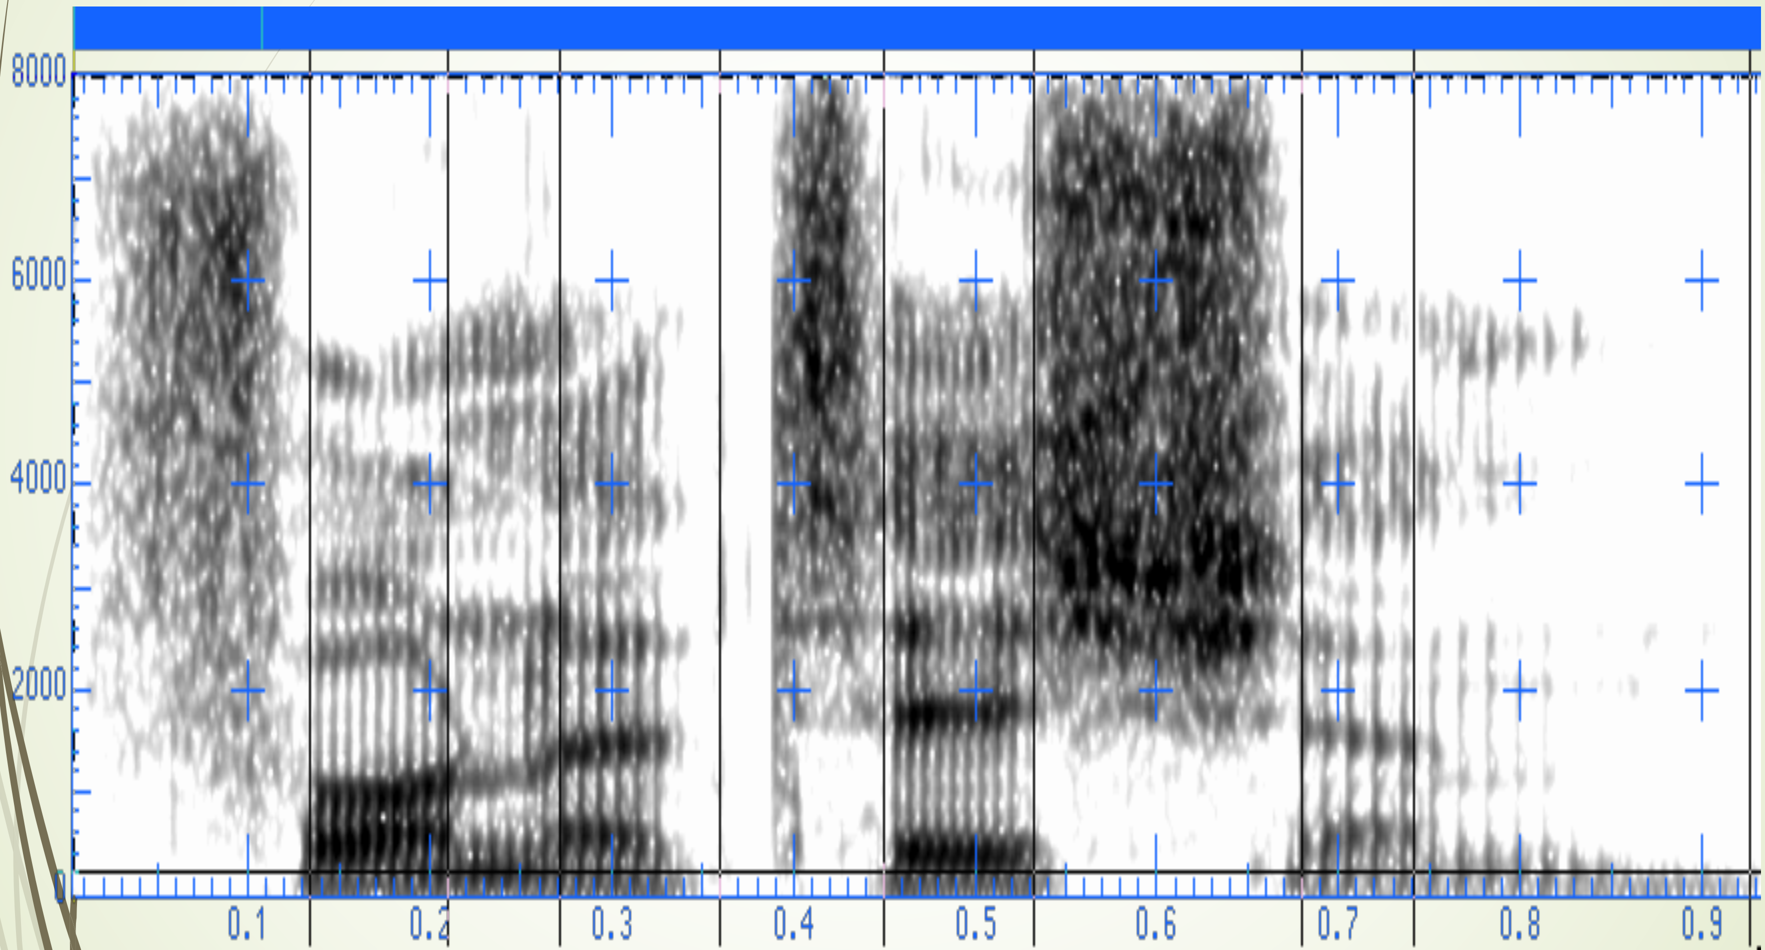

Frequency in cycles per second

Seconds

“Shallow” natural-language processing

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

- Recognize spoken words from a small set of possibilities, e.g. 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 "powerboat" as the object having that data. We can enter this into a ship database.
- This also suggests associations between "Chinese" and "powerboat", and between "powerboat" and "leaking".

Spanish translations from Google Translate

Google Translate builds translations from pieces. Translations of pieces are found by matching documents to their identified translations.

1. "network security" -> "seguridad de red"
2. "software suite" -> "paquete de programas"
3. Combining 1 and 2: "Network-security software suite" -> "paquete de software de seguridad de red"
4. "Purchase our new software" -> "Compre nuestro nuevo software"
5. Combining 3 and 4: "Purchase our new network-security software suite" -> "Compre nuestro nuevo paquete de software de seguridad de red"

Reverse translations may differ since translations are not unique.

Deep 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.
- Uses linguistic theory.
- Needed when the exact content and context matters, as in legal reasoning.
- Deep methods will eventually surpass shallow methods since we have much room to improve deep methods but not shallow methods. Languages change only slowly so they are a near-fixed target.

Example of a parse tree

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

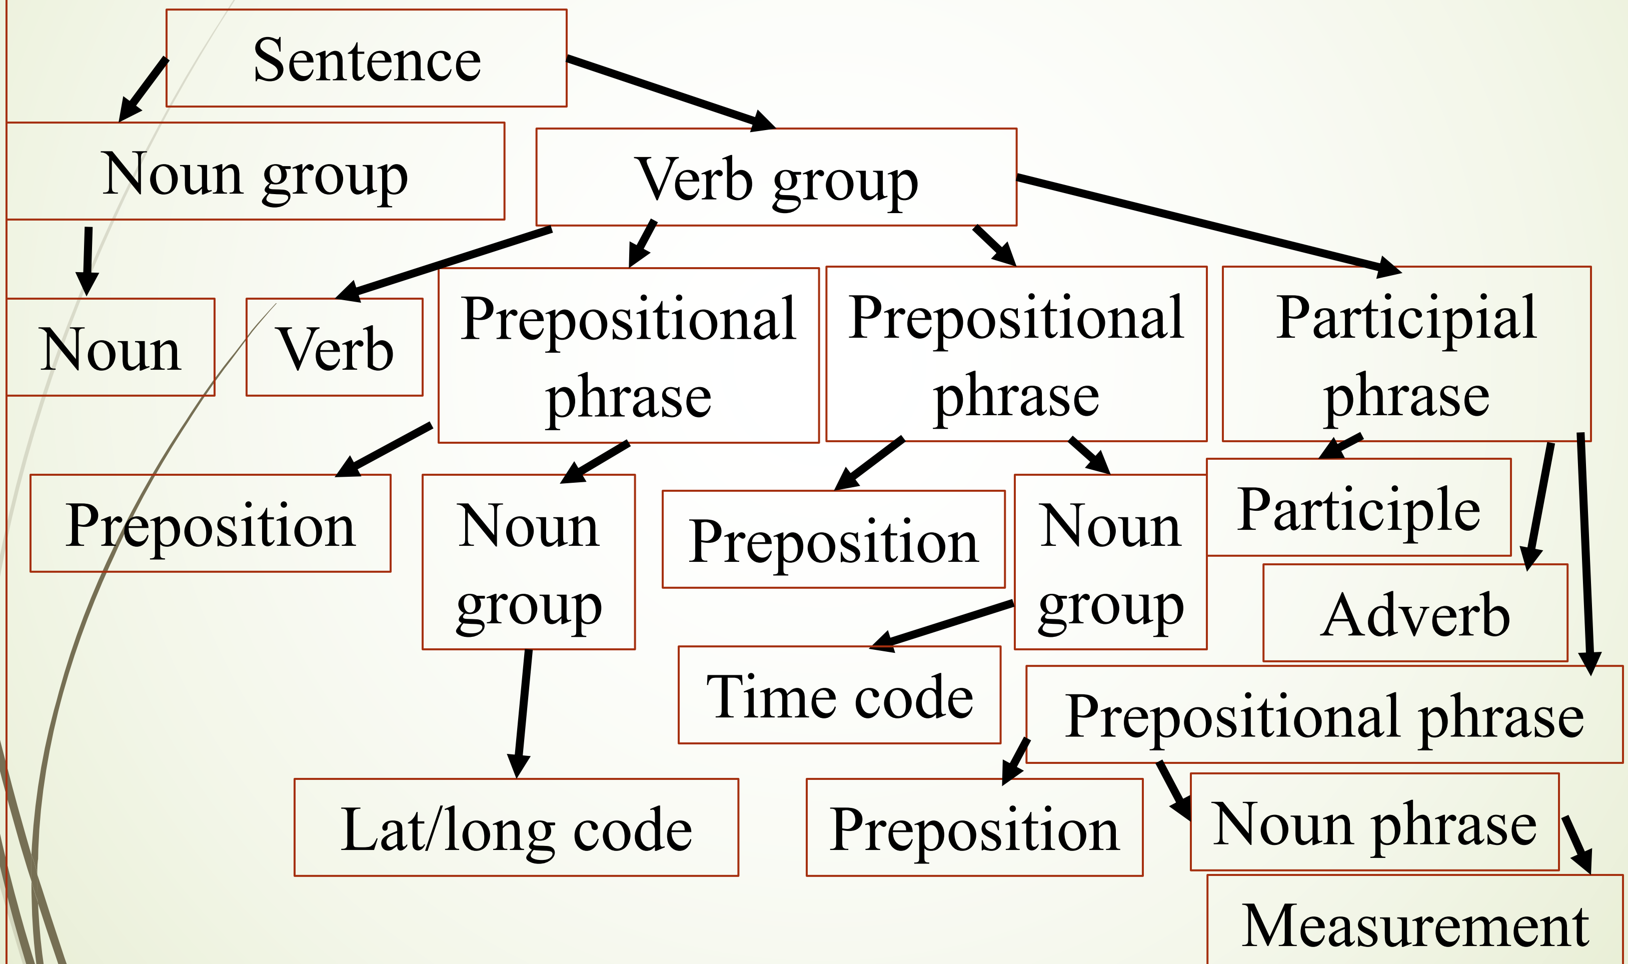

Example of semantic analysis

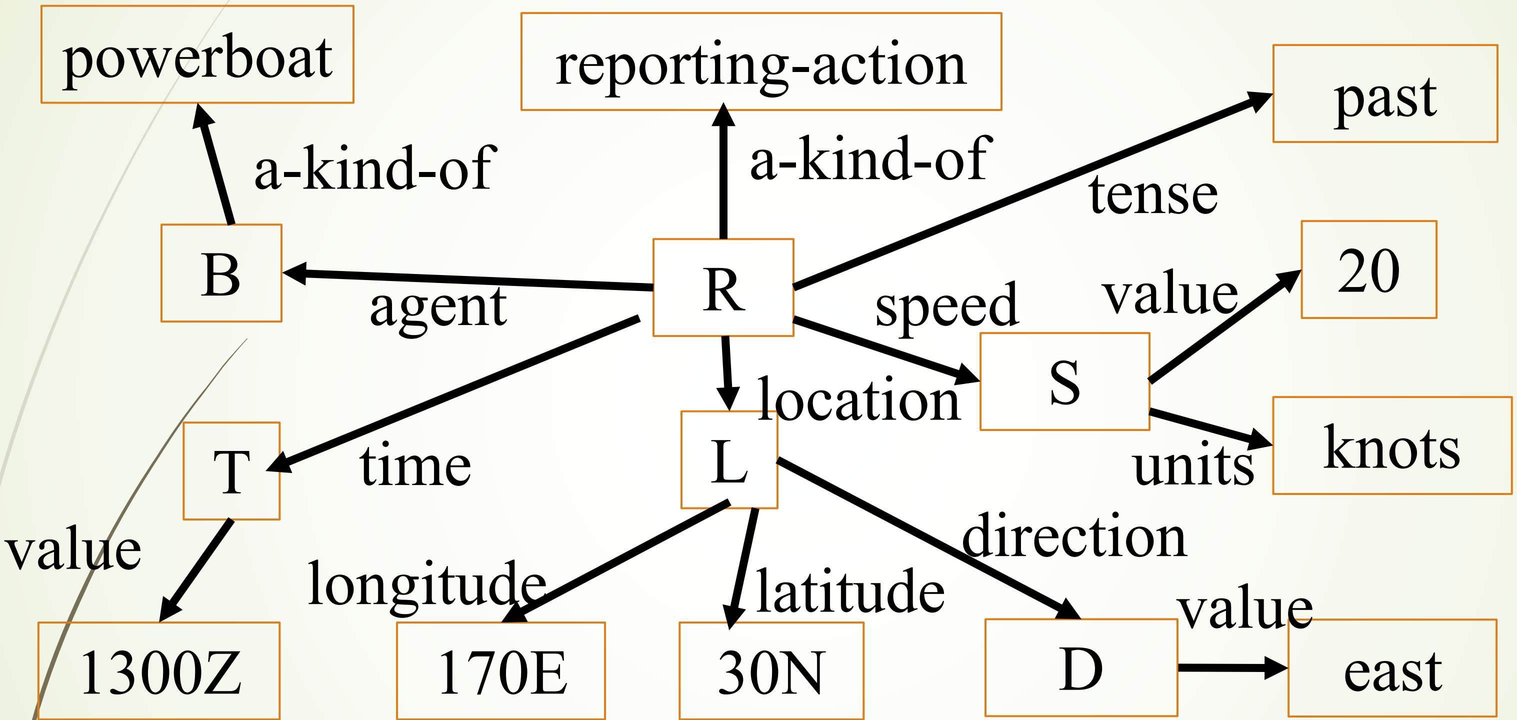

"Powerboat reported at 30N 170E at 1300 heading east at 20 knots.". R is a reporting action, B is a boat, S is a speed, L is a location, T is a time. Processing must figure out these relationships.

Ontology information can help

Unmarked arrows represent type-subtype relationships.

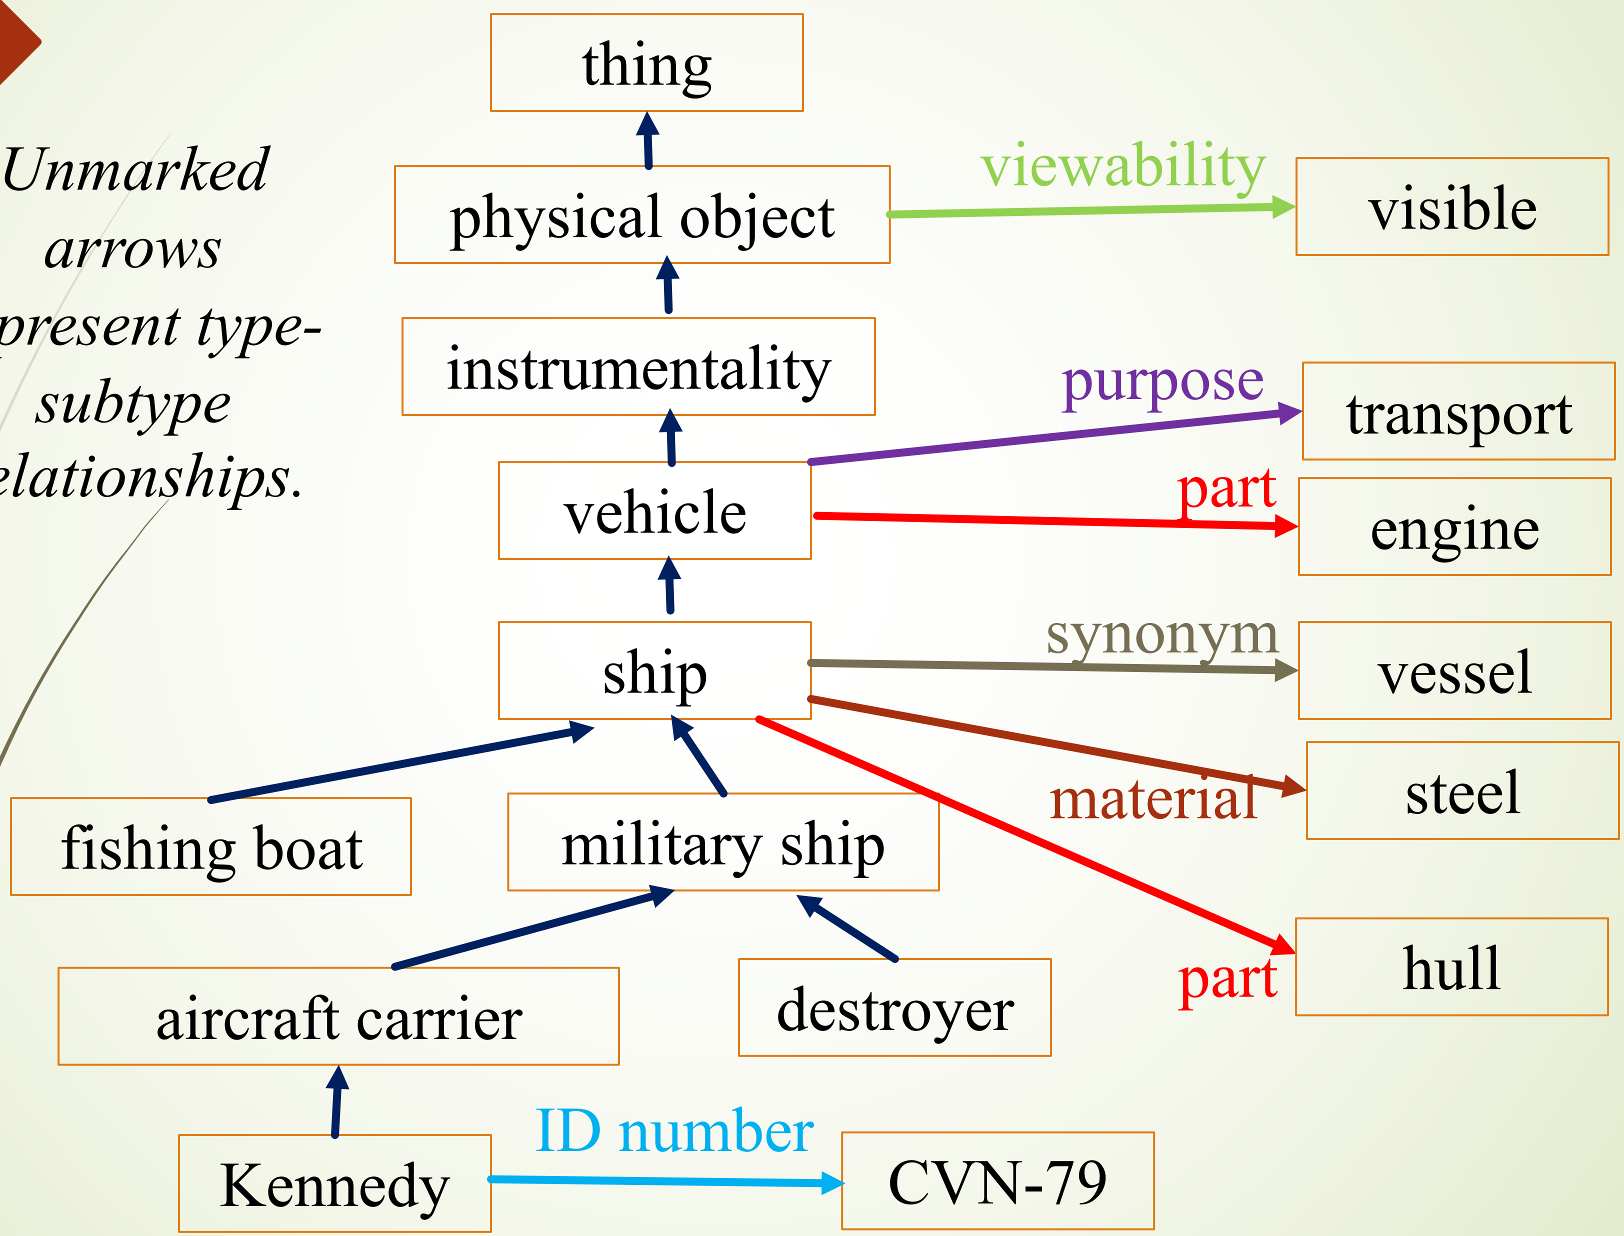

Some properties of money from the FrameNet ontology

Sleep
Sleep_wake_cycle
Smuggling
Soaking
Soaking_up
Sociability
Social_behavior_evaluati
Social_connection
Social_desirability
Social_event
Social_event_collective
Social_event_individuals
Social_interaction_evalu
Socially_significant_hist
Sole_instance
Sound_level
Sound_movement
Sounds
Source_of_getting
Source_path_goal
Spatial_co-location
Spatial_contact
Speak_on_topic
Specific_individual
Speed_description

ascript:openFrame("Sound_movement")

Creator [cre]	The government (or possibly other organization) that mints, issues and guarantees the legal use of the Money.
Inherent_purpose [Inh]	The use for which the Money is intended or designed.
Material [mat]	Any indication of what makes up the Money, including components, ingredients, etc.
Name [nam]	The term used to refer to the Money.
Origin [ori]	The process or means by which the MONEY comes to be possessed. It turns out that the company used drug MONEY to develop this regionv
Possessor [pos]	The Possessor has control over the use the Money. They managed to scam him out of most of his CASH.
Time_of_creation [tim]	The time at which an Money comes into existence or begins to be circulated.
Type [typ]	An indication of the subtype of Money.
Use []	A particular dedicated plan that the Possessor has for how to spend the Money.

From framenet.icsi.berkeley.edu

Example of deeper discourse analysis

- More complex world knowledge is needed to understand reasons for statements and their sequence.
- "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:
  - That the powerboat is in trouble, from knowledge of what "leaking oil" means.
  - That it needs maintenance.
  - More complex world knowledge is needed to understand reasons for statements and their sequence.
  - That it must get maintenance at a port.
  - Leaking oil depletes ship functionality and needs to be fixed soon and nearby.

Miscellaneous problems in natural-language processing

- Understanding also requires:
  - Deciphering cross-references
  - Theories of causation
  - Knowledge of how arguments are constructed
  - Knowledge of social interactions between people ("speech acts")
  - Knowledge of details of common specialized situations (e.g., businesses, law, and politics)
- People demand high accuracy in natural-language conversations: 95% accuracy is not enough.
- Language generation: Not too difficult with a grammar and a semantic representation of what you want to say.