The Orion Project

Summary of Means, Ends, and Issues

General Overview

- The main goal of the project is to build a machine capable of processing human language in written form
- The machine will have an active, undirected, dynamic structure
  - Active: all features will be represented in the form of visible states
  - Undirected: Information can flow through the network in any direction, allowing for backwards or sideways modification
  - Dynamic: The structure can be self-modified to accommodate new information and previous structures can be extended to encompass new inputs
- The machine may be initially tasked with text interpretation in domains such as complex technical specifications, but its utility will not be limited to that

Why?

- An intended tangible benefit of the machine may come in tracking complex logical structures embedded in texts that are difficult for the human mind to follow due to limitations such as memory, or impractical for quick usage due to the time investment required to search through all relevant texts
- Error identification within large texts can be handled by such a machine as it can more readily track the internal logic of a text to identify inconsistencies or violations that may be repeatedly overlooked by human readers

How?

- Firstly, the machine must be able to interpret individual statements (a huge task in and of itself) and weigh the meaning of those statements in relational to subsequent information
- Accomplishing this requires active representation of the logical structure of the analyzed text. This means not only retaining the meaning of statements that have been encountered in the text, but tracking the relationship of those statements to other statements using operators such as AND, OR, NOT, IF, WHEN, etc.

Language Processing Requirements

- Machine must store a large number of word entries and access them at the appropriate time
- In some cases, accessing the appropriate word meaning necessitates deciding between dozens of different entries
  - A means of determining which entry is an appropriate fit for the sentence is required
- Must take into account contextual information when interpreting text
  - This is the primary means of resolving ambiguity in word interpretation. In some cases, interpretation ambiguity may need to be resolved at a later point by incorporating contextual information subsequently encountered within the text
- Must be able to dynamically modify its structure to deal with novel inputs

Dictionary Structure

- We’re currently using a shallow entry-sense-subsense structure built using the OED API
- Problems implementing this structure within our machine:
  - The OED was designed to aid human interpretation, which does not always readily translate into accurate interpretation when used by a machine for whole-text interpretation
  - In some cases there are a very large number of senses/subsenses built into a definition which necessitates a fine-tuned means of definition navigation
  - The sense/subsense relations within the OED, though intuitive to humans, do not always map consistently onto a machine-readable organizational structure
  - This structure may not mirror the way that the human mind stores and accesses information, which has been optimized for language interpretation
- Some of these definitions will need to be reorganized by humans before they can be readily used in machine applications

Alternative Structure

- A natural structure may be a more intuitive means of arranging entry definition storage
- Rather than have the artificial limitation of three-level depth imposed by dictionary structure (entry-sense-subsense), entries would extend as many levels deep as needed:
    - This likely more closely resembles the structure of the human mind
    - If we are to continue building off the OED API, switching to this new form of organization would require heavy restructuring of the accompanying definitions
- The alternative to reorganizing the OED definitions in this way is modify their present structure into a machine compatible form, but this form may not be as well optimized for language interpretation as the human mind

Definition Identification

- We’ve broken down the subsenses into transitive/intransitive/etc. by using the surrounding words to categorize the word in question and simplify its identification
  - For this method to work seamlessly, at least some of the surrounding words need to have clear definitions which may be a problem if the surrounding words also require identification of surrounding words to establish their own definition
- This breakdown may work for verbs, but other types of words will need to use different criteria
  - Prepositions may use the type of relationship they convey (temporal, spatial, directional, etc.)
  - Literal/figurative division may be a productive place to start for nouns

Resolving Ambiguity

- A single word may have dozens of different meaning entries. How does one decide which entry is the appropriate one?
- Context
  - Themes within the text, paragraph, or sentence may provide clues as to which is the intended meaning of a word
- Syntactic cues
  - Word order and other syntactic rules may provide cues to what role a word plays in the sentence (noun, verb, etc.), which may help in selecting the appropriate definition
- Selecting an appropriate meaning results from forming prediction models that bias us towards certain interpretations of the sentence, carrying out process of elimination, and using our “best guess” (essentially probabilistic) when ambiguity cannot be fully resolved and waiting for more information is not possible

Integrating Context Information

- In the human mind:
  - Language processing occurs “online”. As any given utterance is being constructed, the mind makes predictions about likely information that will follow
  - These predictions bias the interpretation of subsequent word choices. These prediction models take into account base word usage rates/contextual information and serve to resolve any ambiguity in the utterance as well as to speed interpretation
  - As further information becomes available, previous “best guess” interpretations may be overwritten to be brought into line with that information
- Any machine capable of interpreting language in the way that humans do needs to be able to predict the most likely meanings of words, and track the fit of those meanings with subsequently encountered information, adjusting the definition choice as necessary

Integrating Context Information

- Machine Approaches:
  - Connectionist models have been used to represent language as a series of connected nodes within a “neural network”
  - These nodes have particular weights that bias interpretation in particular ways. These weightings can be dynamically adjusted based on context information
  - Though the adjustable weightings of such networks grants them some flexibility, their framework precludes reorganization of existing structures to accommodate new information (knowledge transfer)
- Our machine will favor an active structure so as to avoid the pitfalls of typical neural network approaches

Active Structure

- No hidden aspects; the structure is made up entirely of visible states rather than relying upon non-visible instructional sets
  - This means that instructional operations normally represented in underlying algorithms will need to be translated into an active state form
- Structure can be self-modified to handle new inputs
  - This self-modification is a reason why an algorithm-dependent instantiation should not be used: the structure that an underlying algorithm/instructional set depended on being there may change
- Using an active structure facilitates integration of preexisting relationships into new representational structures when encountering novel situations

Language Usage Issues

- There are a variety of language-related issues that will need to be solved over the course of development such as:
- Idiom usage
  - Idioms often require heavy abstraction of meaning and application in an unconventional way. Will idiomatic usages need to have discrete entries, and how will the machine deal with novel forms it may encounter?
- Resolving ambiguities in rule usage
  - In “He wanted to call her name” vs “He wanted to call her names”, why does our interpretation shift away from the possessive form of ”her” with the pluralization of “name”?
- Establishing pronoun referent
  - In addition to their singular form, pronouns such as ”they” and ”them” may represent any number of subjects. The singular vs plural form can be established by subject agreement within the sentence, but in the case of the plural form how will a machine know whether to take the previous two agents in the text rather than the previous three?

Dynamic Self-Modification

- Given the sheer quantity of possible utterance combinations and the ever-changing nature of language, it is not plausible to build a machine capable of handling all utterances while relying upon static predetermined structure
- Any competent language processing machine needs to be able to adapt its own structure dynamically and extend the logic it has used in previous instances to new cases

What to prioritize first?

- There are obviously many complex steps that need to be completed in this project but the steps we need to first address are:
  - Determine whether we wish to stick with the entry-sense-subsense organization or reorganize into a natural structure
  - Clean up dictionary definitions into a consistent and machine-useable form
  - Fine-tune the relational structures between the existing dictionary definitions

Active Structure