
Proceedings Paper:
Atwell, ES (1987) Natural language computing in the office. In: McManus, B and Salenieks, P, (eds.) Computers in tomorrow's society: selected papers from the 2nd BCS Young Professionals Group Conference. The 2nd BCS Young Professionals Group Conference, 01 Jun 1987 British Computer Society , 34 - 40.

Natural language computing in the office
Eric Atwell

What is natural language computing?

Systems capable of processing English, or some other natural language such as Japanese or Esperanto, are highly complex, but generally include (most of) the following four components:

- With speech input, the speech waveform has to be converted to ASCII character sequences. If the input is a single word, chosen from a small set of allowable words, then the speech signal can be matched against a set of templates, and the character-string associated with the nearest match is selected. Fig. 1 shows a set of templates for an input word to be matched against. If the input utterance is more complex, syllable and word boundaries have to be worked out first, and this is very complicated. Because of this difficulty, most current systems expect the English to be typed using a keyboard, thus avoiding the sound-to-character conversion problem.

- Each word in the input has to be looked up in a dictionary, to find associated information needed in subsequent processing, such as grammatical parts of speech (noun/verb/adjective etc), meaning, etc. In advanced systems, each dictionary entry may contain a complex structure of information. Fig. 2 illustrates the structure of entries in LDOCE, a dictionary used by many researchers including those working on IBM and Alvey projects.

- A parser then works out the grammatical structure of the sentence, for example, dividing it into subject, verb and object(s), and dividing each of these phrases in turn into smaller constituents. The parser produces a tree structure such as the one shown in Fig. 3.

- Such a syntax tree is then used to work out the semantic or meaning representation of the sentence. Just what this looks like depends very much on the application. For some applications, the semantic representation must be very detailed, with complex structures even for apparently simple sentences; for example, Fig. 4 shows the analysis given by one system for the sentence "John shut the door."

Many current systems for natural language processing were designed by academics in artificial intelligence (AI) and computational linguistics, to test specific theories of how humans process language. As in other areas of AI, academic prototypes are giving rise to commercial applications. In principle, applications are only limited by human imagination - with possible areas of benefit as far apart as educational and military (including espionage and surveillance). The office is the biggest area currently targeted.

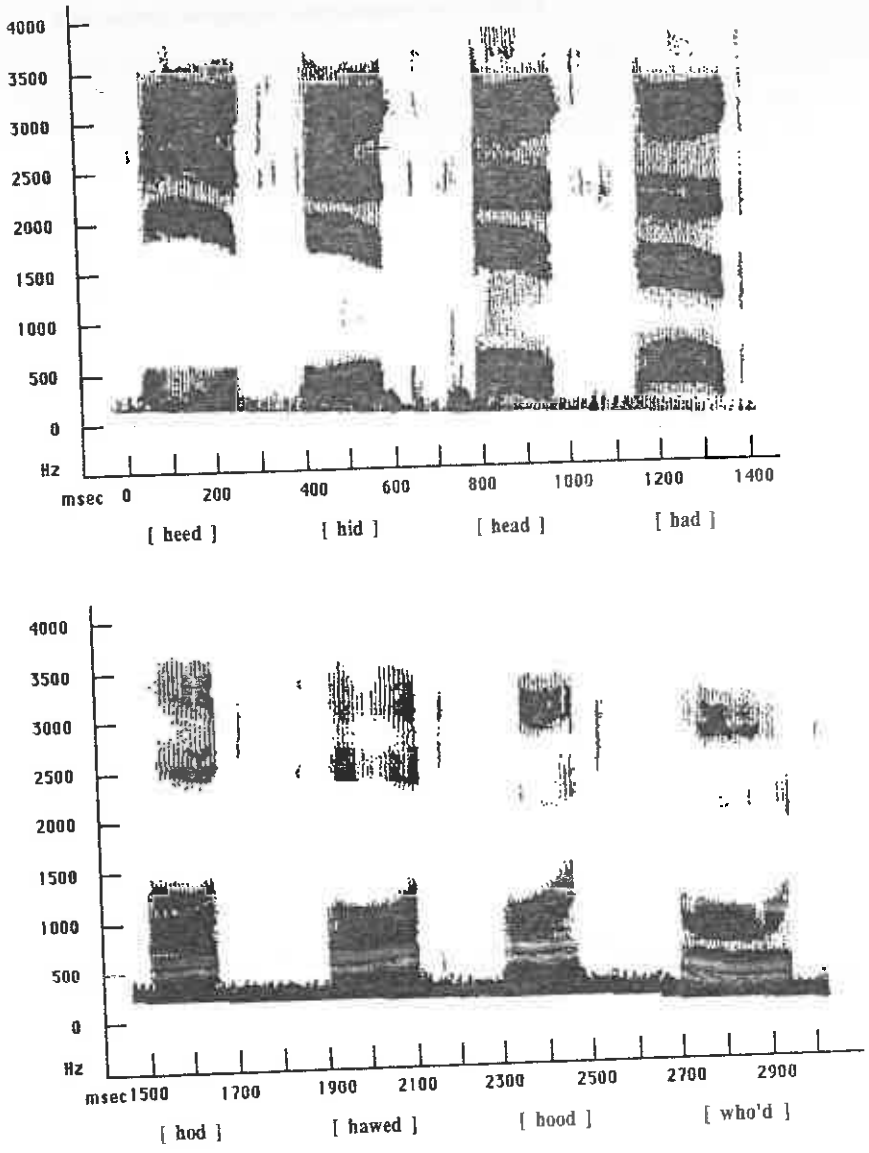

Figure 1 A spectrogram of the words: heed, hid, head, had, hod, hawed, hood, who'd

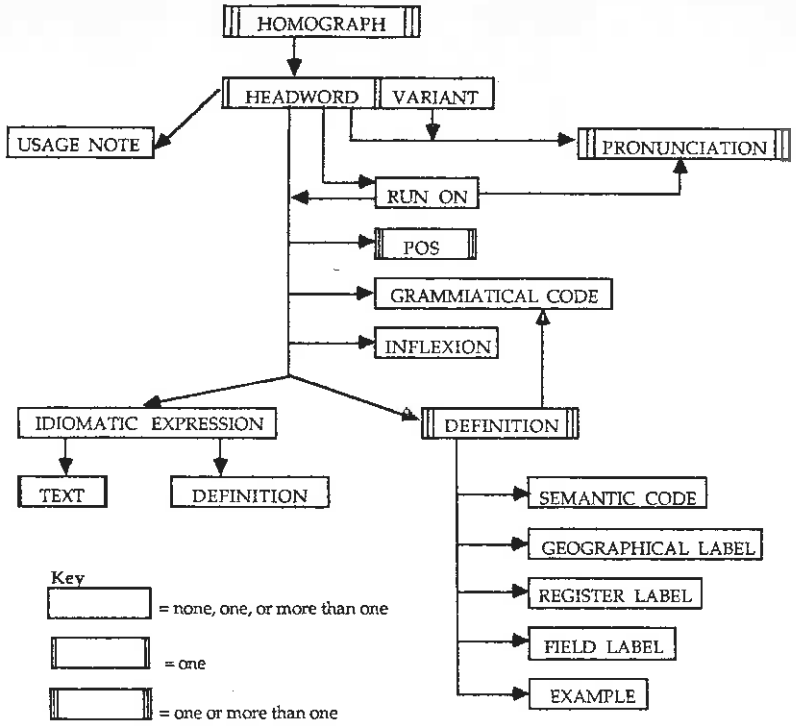

Structure of a compound homograph

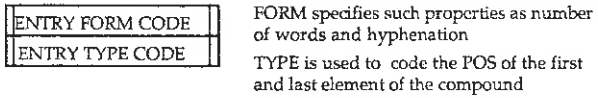

FORM specifies such properties as number of words and hyphenation

TYPE is used to code the POS of the first and last element of the compound

Figure 2 Structure of an LDOCE entry

Why natural language? - applications in the office

Office applications of natural language computing divide into three main areas:

- Single-word speech input - e.g. for driving menu-based interfaces to software systems. A word spoken by the user is matched against a set of templates, and the best match chosen; so the utterance could be a phrase instead of a single word.
- Interfaces which require fairly complex commands to be given via a keyboard. Other user-friendly input techniques (menus, pointing with a mouse, etc) are too cumbersome and slow; but formal command or query languages have to be learnt before they can be used. Natural language interfaces allow users to issue commands or queries in English, and so can be used by the casual or novice user without training. Examples include database query systems, expert system consultation and interrogation. Commercial systems still require English to be typed in via a keyboard; current research aims to take speech input beyond isolated-word recognition to deal with continuous speech, that is, complete sentences.
- Currently, a great deal of typed English (or other natural language) text has to be analysed or processed by humans in some way. Natural language computing techniques will allow much of this to be automated. Example tasks include: machine translation, automatic content-scanning and abstracting, proof reading word-processed documents. Again, current systems expect typed text input. IBM, Kurzweil, Alvey, etc are funding major research efforts to automate the typist too, by developing automatic dictation-transcription systems (or talkwriters).

Commercial natural language computing systems have only recently reached the marketplace. A recent survey estimated that total UK sales in 1985 were about three quarters of a million pounds, and in the USA about ten million pounds. However, as more academics open their expertise to commercial exploitation, natural language computing should become increasingly commonplace in the office. The same survey predicted that in ten years, sales would mushroom to over 150 million pounds in the UK, and over one billion pounds in the USA! Clearly, natural language computing will be one of the key features in computers in tomorrow's society.

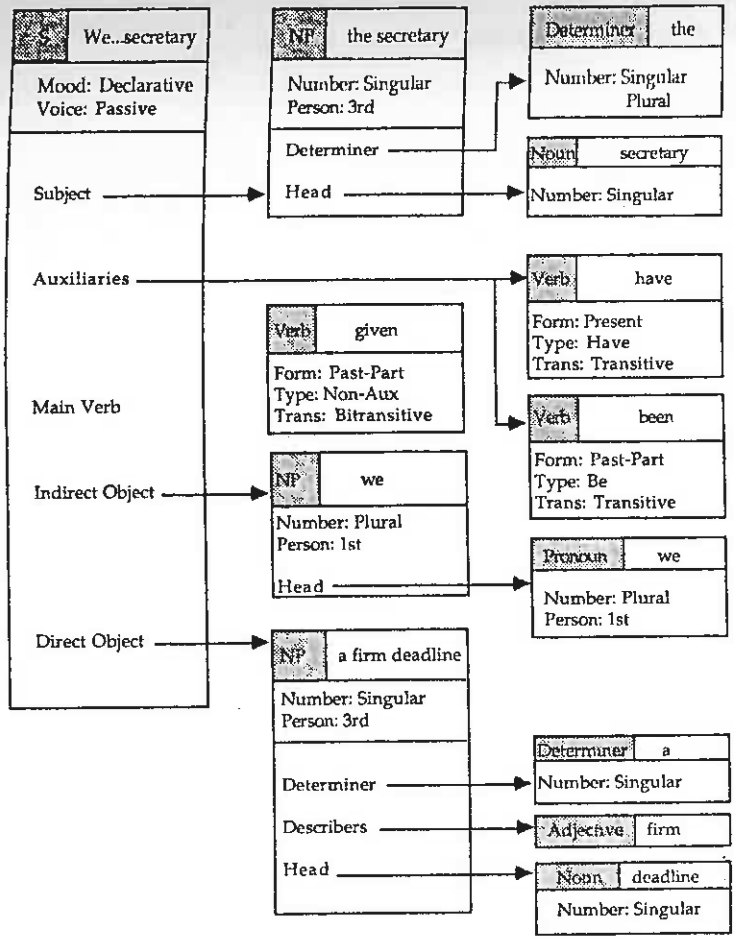

Figure 3

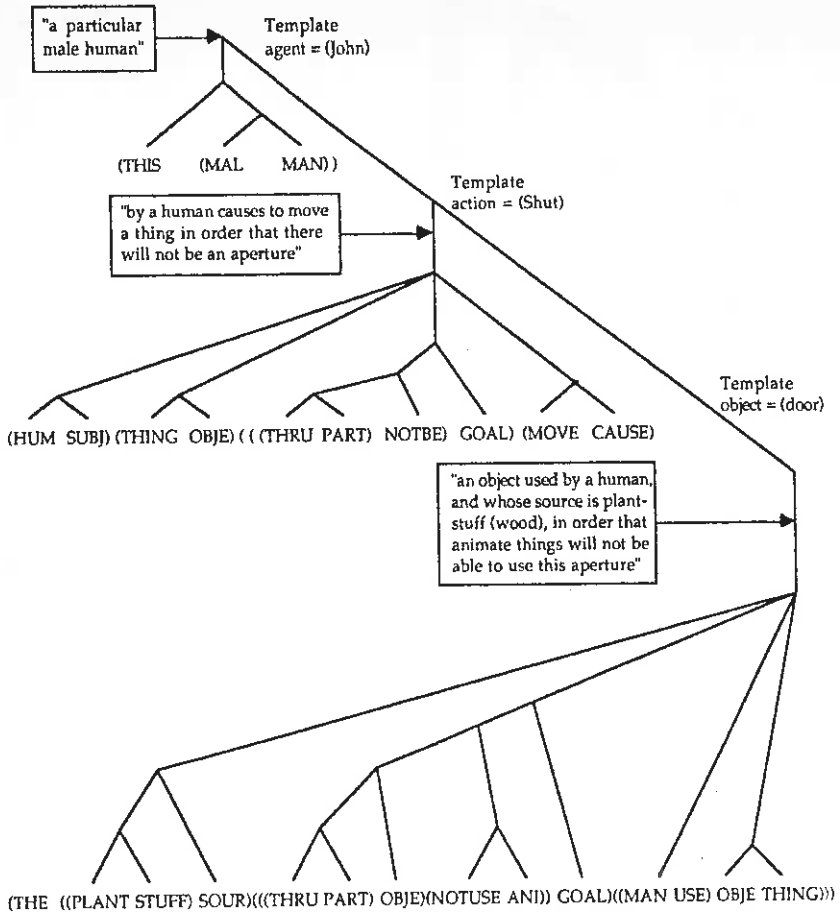

Figure 4