A Computational Model Of
Natural Language Understanding Based On
Auditory Neuroscience and Graph Theory

Lloyd Watts
July 31, 2022

Abstract – Recent advances in Auditory Neuroscience have shown the cortical brain regions responsible for decoding spoken language and organizing it into hierarchical structures, using previously stored knowledge of Entities and Events. This paper describes an experimental computational model for Natural Language Understanding based on those structures, using a structured dictionary to represent the previous stored Entity and Event knowledge. Given an input sentence, the model attempts to find a coherent understanding of the meaning of the sentence, out of the many possible syntactical and semantic combinations of the individual word meanings. The most coherent understanding is the one which maximizes a score based on the sum of the inverse distances between word pairs in the dictionary word graph, as determined by a bidirectional breadth-first search. With a well-designed dictionary, this approach finds good understandings of the sentences, and can justify its understanding by showing the relationships between the chosen senses of the words using a backtrace of the best paths.

1. Introduction

Modern Natural Language Processing systems based on Deep Learning and Transformers have achieved State of the Art performance on various language tasks, i.e., over 80% accuracy on Word Sense Disambiguation [1][2]. Modern language-generating systems based on Deep Learning and Transformers such as OpenAI's GPT-3 are capable of completing sentences in ways that are sometimes impressively human-like. However, these systems have been criticized as simply finding patterns in their training data, and not really understanding the meaning of their inputs or their outputs [3]. From Gary Marcus:

> "At first glance, GPT-3 seems to have an impressive ability to produce human-like text. And we don't doubt that it can be used to produce entertaining surrealist fiction; other commercial applications may emerge as well. But accuracy is not its strong point. If you dig deeper, you discover that something's amiss: although its output is grammatical, and even impressively idiomatic, its comprehension of the world is often seriously off, which means you can never really trust what it says." [4]

Another criticism from Walid Saba is that humans communicate with language, relying on a shared understanding of common sense concepts, such that the communicated words are not sufficient on their own to understand the meaning of the message.

> "When we communicate, we leave a lot of things out. So for example, I say, 'Mary enjoyed the movie.' I didn't say watching is what she enjoyed, but obviously what she enjoyed is watching,

> although people can direct and produce and sell and buy movies, but you know I meant Mary enjoyed watching them. Or Mary enjoyed the book, meaning Mary enjoyed reading the book. This is just a small example of what I call in my writing 'the missing text phenomenon.' The text that we read or we hear is just a clue to the thought I'm trying to convey. We leave a lot of stuff out because we all have common sense knowledge. We know, we all know what we all know, so the data itself doesn't have all the information. No matter how much you analyze it, it's missing something else to do the full comprehension. Without getting too technical or into too much detail, the data itself is one part of the puzzle. The other part is common knowledge that has to be engineered in the system to fill in the gaps. Because machines don't know what we all know. That's why it's difficult for machines. So we need approaches that can discover the missing text, the missing information, that's not in the data. Which makes it very challenging." [5]

Another criticism of modern Large Language Model systems is that they are not grounded in Neurobiology [6]. They are purely computational and statistical models that are not informed by what we know about the architecture of the human brain, and how the human brain works to comprehend language.

Recent advances in Auditory Neuroscience have shown the cortical brain regions responsible for decoding spoken language and organizing it into hierarchical structures, using previously stored knowledge of Entities and Events. From Greg Hickok: "In comprehension, the posterior middle temporal gyrus (pMTG) functions to decode sequences of auditory phonological representations in the posterior superior temporal gyrus (pSTG) into hierarchical structures and link these with two conceptual networks, an entity knowledge hub in the anterior temporal lobe (ATL) and an event knowledge hub in the angular gyrus (AG)." [7][8] These brain regions are shown in Figures 1 and 2. Taken together, these regions comprise the Natural Language Understanding system in the human brain.

Could we build a Natural Language Understand system that addresses the above criticisms, and builds on what we know about how the human brain processes language? It would be insufficient to rely entirely on an unlabeled corpus of text. It would have to have some explicit representation of the meaning of words, and the relationships between the words in their various possible meanings, and demonstrate a coherent understanding of the meaning of sentences it receives as inputs.

The objective of the present work was to build a computational model of this system with certain properties:

- It should take English sentences as inputs, and produce an annotated output sentence showing the best meaning of the sentence, including the best choice of Part of Speech (Parsing) and Word Meaning (Word Sense Disambiguation) for each word in the sentence.
- It may have access to a suitably structured English-language dictionary, representing the accumulated human adult knowledge of English vocabulary and word definitions.
- It should be Explainable, i.e. it should be able to justify its selection of best sentence meaning, and show its best choice of relationships between the words.
- It should operate in a way that is consistent with what we know about language processing in the human brain [7].
- A Proof of Concept could be built in a few days by a single engineer, as a feasibility study to precede a larger team effort.

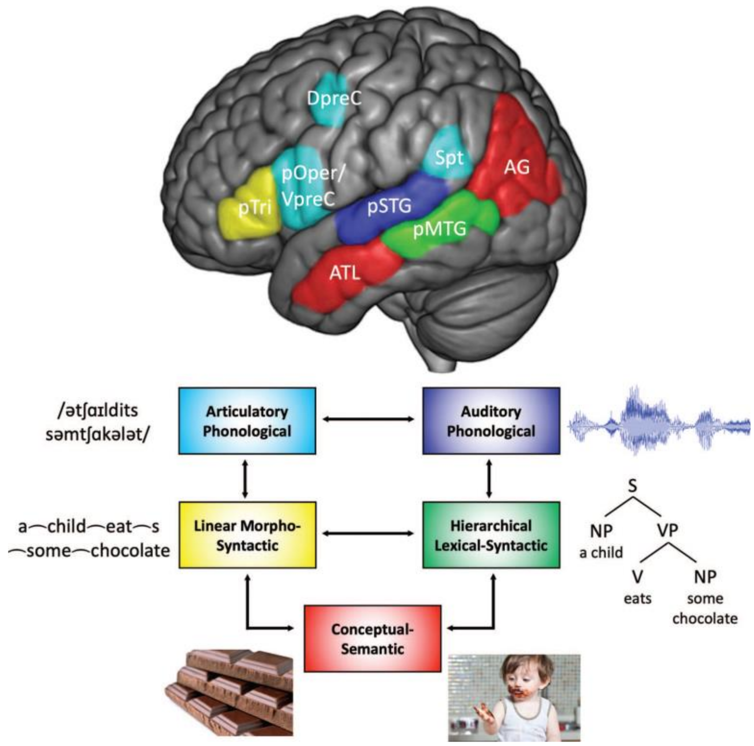

Figure 1: Cortical organization of syntax and the sensorimotor and conceptual-semantic brain systems it interfaces with [7].

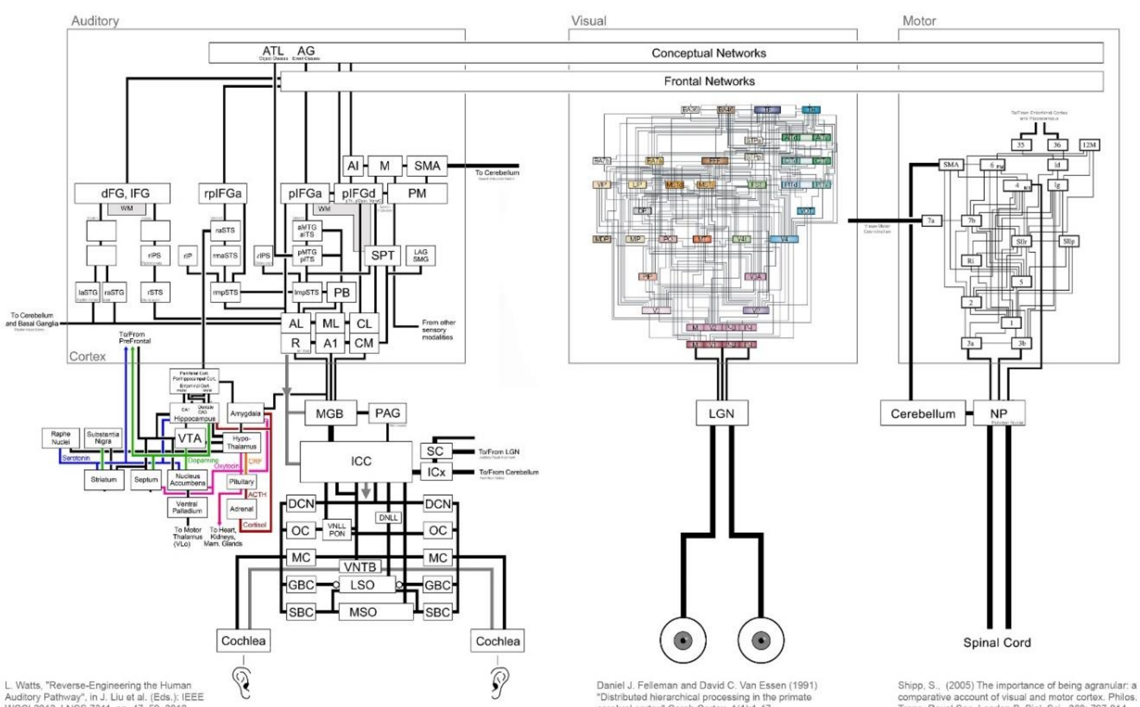

Figure 2: Architecture of the Human Brain, showing the Anterior Temporal Lobe (ATL) with object classes, and Angular Gyrus (AG) with event classes, in the Conceptual Networks associated with the Auditory Pathway [6][9].

Dictionary-based methods have been used in the past. The first usage was by Michael Lesk in 1986, who used a simple word-pair distance metric based on the number of common words in their definitions [10]. Further extensions of these methods are surveyed by Roberto Navigli in 2009 [11]. These methods fell out of favor in the 2010's with the rapid adoption of Deep Learning and Transformer methods.

The present work is a Proof of Concept, intended to show that the approach can work on some interesting example sentences, and to demonstrate the form of its outputs and behavior on those examples. But full scaling, benchmarking, robustness testing, and performance optimization is beyond the scope of this initial experiment.

2. A Computational Model of Natural Language Understanding

Figure 3 shows a block diagram of our computational model of Natural Language Understanding.

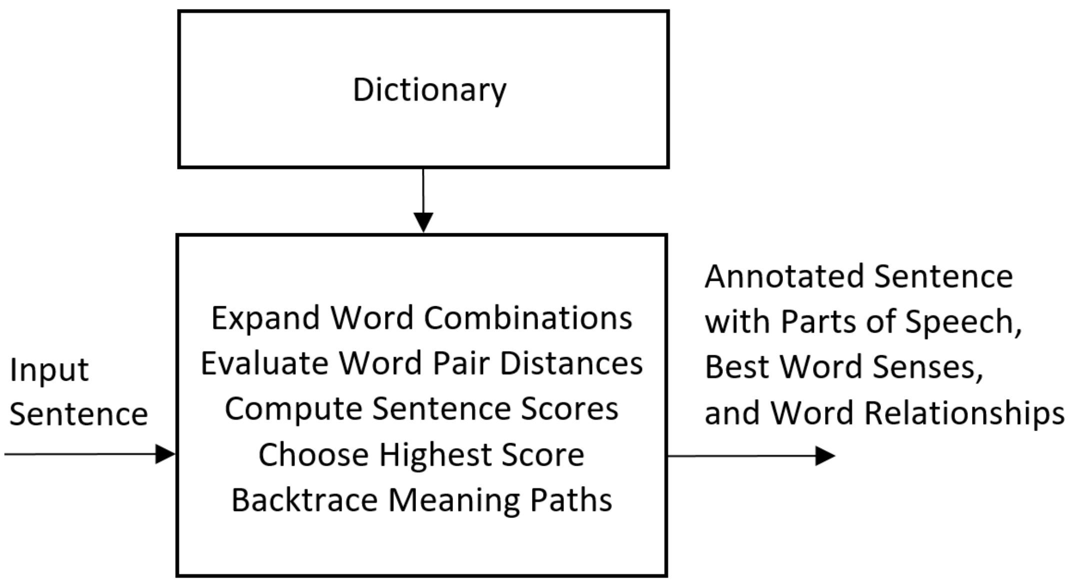

Figure 3: Computational Model of Natural Language Understanding.

To describe the approach, we will walk through a simple example, for the input sentence "The cat sat on the mat." In the first step, the model considers the possible meanings of each word. For example, the word "cat" has a total of 9 meanings in the dictionary, 7 nouns (NN) and 2 verbs (VB). The word "sat" is the past tense of the verb "sit", with a total of 13 possible meanings in the dictionary.

```
sentence = The cat sat on the mat.
words = The cat sat on  the mat .

Word           Meanings  Possible Parts of Speech        Meanings
----           --------  ------------------------        --------
The                   1  ['DT']                          [1]
cat                   9  ['NN', 'VB']                    [7 2]
sat                  13  ['VBD']                         [13]
on                    8  ['JJ', 'RB', 'IN']              [2 3 3]
the                   1  ['DT']                          [1]
mat                   9  ['JJ', 'NN', 'VB']              [1 6 2]
.                     1  ['.']                           [1]
```

Considering all the possible combinations of meanings, there are 9*13*8*9 = 8424 possible sentence meanings. And considering all the possible combinations of Parts of Speech, there are a total of 2*3*3 = 18 possible sentence types.

```
number of raw meanings = 8424
number of sentence types = 18
```

The table below shows a listing of the 18 possible sentence types, with the number of meanings associated with each sentence type. Some of these sentence types can be eliminated from consideration on the basis of some simple and obvious syntax rules. For example, the second-last word in the sentence is "the", so therefore the last word must be a noun (NN). Thus, we can mark any of the sentence types with a last word used as a verb (VB) or adjective (JJ) as not allowed. Applying a few simple rules eliminates all but one sentence type in this example. So the model has reduced the number of combinations to consider from 8424 to 1638.

```
All Possible Sentence Types:     Meanings   Allowed
----------------------------     --------   -------
The cat sat on  the mat .        8424

DT  NN  VBD RB  DT  NN  .        1638       False
DT  NN  VBD IN  DT  NN  .        1638       True
DT  NN  VBD JJ  DT  NN  .        1092       False
DT  NN  VBD RB  DT  VB  .         546       False
DT  NN  VBD IN  DT  VB  .         546       False
DT  VB  VBD RB  DT  NN  .         468       False
DT  VB  VBD IN  DT  NN  .         468       False
DT  NN  VBD JJ  DT  VB  .         364       False
DT  VB  VBD JJ  DT  NN  .         312       False
DT  NN  VBD RB  DT  JJ  .         273       False
DT  NN  VBD IN  DT  JJ  .         273       False
DT  NN  VBD JJ  DT  JJ  .         182       False
DT  VB  VBD RB  DT  VB  .         156       False
DT  VB  VBD IN  DT  VB  .         156       False
DT  VB  VBD JJ  DT  VB  .         104       False
DT  VB  VBD RB  DT  JJ  .          78       False
DT  VB  VBD IN  DT  JJ  .          78       False
DT  VB  VBD JJ  DT  JJ  .          52       False

number of allowed sentence types = 1
number of allowed meanings = 1638
```

So the model has reduced the number of combinations to consider from 8424 to 1638. Now, for each of those 1638 combinations of word meanings, we will compute a graph theory distance between each of the nontrivial word pairs. In this sentence, there are 4 nontrivial words: cat, sat, on, mat. There are a total of C(4,2) = 6 word pairs: cat/sat, cat/on, cat/mat, sat/on, sat/mat, and on/mat. Let's consider the graph theory distance calculation for the correct meanings of cat and sat:

- "cat": "def": "feline mammal usually having thick soft fur and no ability to roar, usually domesticated",
- "sat: "def": "past tense of sit"; "sit": "def": "to rest on the buttocks or haunches"

We need to compute the graph theory distance between those two meanings. We will use bidirectional bread-first search [12], in which we recursively expand the definitions of each word until we find a common nontrivial word, and then count the number of levels that needed to be traversed to find the common word. In this case, there are no common words at the first level of definitions. So, we go up one level from "cat", and up one level from "sat", and find no common words, so the distance between them must be greater than 1+1=2. If we go up another two levels from "sat", we get the following sequence of expanded definitions:

- "sat: "def": "past tense of sit"; "sit": "def": "to rest on the buttocks or haunches"
- "haunches": "def": "plural of haunch"; "haunch": "def": "the loin and leg of a quadruped",

- "quadruped": "def": "an animal especially a mammal having four limbs specialized for walking"

And now we have found a common word, "mammal", which is up one level from "cat" and up three levels from "sat", for a total distance of 4, as shown in Figure 4. And we can trace the path between these two definitions as cat -> mammal -> quadruped -> haunches -> sat. This provides a sensible linkage between the words, the hidden common-sense missing information between the words that is present in the overall context of the whole structured dictionary. "Cats can sit because they are mammals with legs."

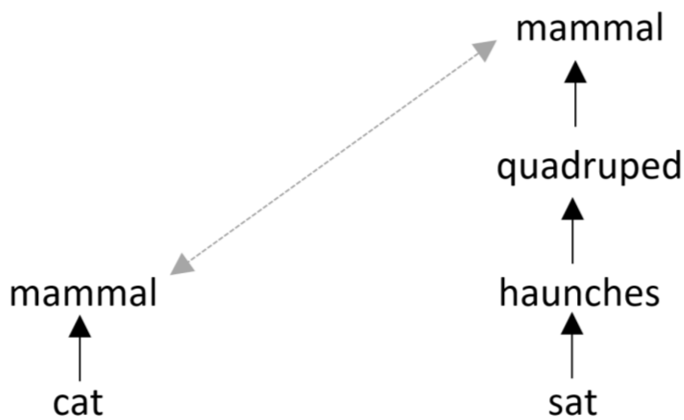

Figure 4: Computing the distance between "cat" and "sat" for the above choice of meanings. The total number of solid upward arrows is 4, to find the common word "mammal" in their recursive definitions. The backtrace path between these two definitions is cat -> mammal -> quadruped -> haunches -> sat.

Notice that the model found the common-sense connection between cat and sat from first principles just using the dictionary definitions and the distance measure. It was not explicitly provided by human annotators, as is done in WordNet [13].

Notice also that for any of the other definitions of "cat", such as "whip (cat-o-nine-tails)" or "boat (catamaran)", or definitions of "sat", such as "pose for a portrait" or "take an exam", the distance to the word "sat" is larger, at least 5.

We will compute the graph theory distance measure for every word meaning pair in a given sentence meaning combination. The score for the sentence meaning combination will be the sum of the inverse distances of all of the word meaning pairs. And then we will choose the sentence meaning combination with the highest score. For the sentence "The cat sat on the mat", the best score was 1.583, for the correct interpretation of the meaning of the sentence, as shown below, including the word definitions and parts of speech.

```
Top 3 Meanings out of 1638 calculated:
sortedMeaningsAndScores[  0] = ['DT', 'NN', 'VBD', 'IN', 'DT', 'NN', '.']  [0, 0, 0, 0, 0, 0, 0]  1.583
sortedMeaningsAndScores[  1] = ['DT', 'NN', 'VBD', 'IN', 'DT', 'NN', '.']  [0, 0, 1, 0, 0, 0, 0]  1.528
sortedMeaningsAndScores[  2] = ['DT', 'NN', 'VBD', 'IN', 'DT', 'NN', '.']  [0, 1, 1, 0, 0, 0, 0]  1.474

Highest Scoring Meaning:
meaningIndices = [0, 0, 0, 0, 0, 0, 0]   score = 1.583
The          DT      the definite article, used to specifically identify something
cat          NN      feline mammal usually having fur and no ability to roar, usually domesticated
sat          VBD     (past tense) to rest on the buttocks or haunches   verbBase = sit
on           IN      used as a function word to indicate position in contact with and supported by the top surface of
the          DT      the definite article, used to specifically identify something
mat          NN      a thick flat pad used as a floor covering
.            .       .
```

```
backTrace Records:
cat -> mammal -> quadruped -> haunches -> sat
cat -> fur -> coat -> surface -> on
cat -> roar -> animal -> pad -> mat
sat -> rest -> land -> surface -> on
sat -> haunches -> human -> pad -> mat
on -> surface -> floor -> mat
```

And of course, the model considered 1637 other combinations of meanings which did not score as well, and were thus rejected. And the model shows the 6 backtrace records, as a way of explaining its perceived connections between each of the words in the selected sentence interpretation.

3. Other Examples

Walid Saba gives these interesting examples:

- Mary enjoyed the sandwich. (eating it)
- Mary enjoyed the book. (reading it)
- Mary enjoyed the movie. (watching it)

He points out that understanding the sentence requires information that is not in the sentence. The model naturally finds the missing information as a by-product of finding the best meaning, as shown in the backtrace.

Mary enjoyed the sandwich.

```
Top 3 Meanings out of 5 calculated:
sortedMeaningsAndScores[  0] = ['NN', 'VBD', 'DT', 'NN', '.']  [0, 2, 0, 0, 0]  0.521
sortedMeaningsAndScores[  1] = ['NN', 'VBD', 'DT', 'NN', '.']  [0, 1, 0, 0, 0]  0.506
sortedMeaningsAndScores[  2] = ['NN', 'VBD', 'DT', 'NN', '.']  [0, 3, 0, 0, 0]  0.481

Highest Scoring Meaning:
meaningIndices = [0, 2, 0, 0, 0]   score = 0.521
Mary         NN      a woman
enjoyed      VBD     (past tense) derive or receive pleasure from   verbBase = enjoy
the          DT      the definite article, used to specifically identify something
sandwich     NN      food with two or more slices of bread with a filling between them
.            .       .

backTrace Records:
Mary -> woman -> _female_ -> accept -> receive -> enjoyed
Mary -> woman -> _human_ -> eat -> food -> sandwich
enjoyed -> receive -> _mental_ -> food -> sandwich
```

Mary enjoyed the book.

```
Top 3 Meanings out of 45 calculated:
sortedMeaningsAndScores[  0] = ['NN', 'VBD', 'DT', 'NN', '.']  [0, 2, 0, 0, 0]  0.631
sortedMeaningsAndScores[  1] = ['NN', 'VBD', 'DT', 'NN', '.']  [0, 3, 0, 0, 0]  0.583
sortedMeaningsAndScores[  2] = ['NN', 'VBD', 'DT', 'NN', '.']  [0, 1, 0, 0, 0]  0.556

Highest Scoring Meaning:
meaningIndices = [0, 2, 0, 0, 0]   score = 0.631
Mary         NN      a woman
enjoyed      VBD     (past tense) derive or receive pleasure from   verbBase = enjoy
the          DT      the definite article, used to specifically identify something
book         NN      a written work or composition that has been published in the form of pages bound together to be read
.            .       .

backTrace Records:
Mary -> woman -> _female_ -> accept -> receive -> enjoyed
Mary -> woman -> _human_ -> read -> book
enjoyed -> receive -> register -> _book_
```

Mary enjoyed the movie.

```
Top 3 Meanings out of 5 calculated:
sortedMeaningsAndScores[  0] = ['NN', 'VBD', 'DT', 'NN', '.']  [0, 2, 0, 0, 0]  0.631
sortedMeaningsAndScores[  1] = ['NN', 'VBD', 'DT', 'NN', '.']  [0, 1, 0, 0, 0]  0.597
sortedMeaningsAndScores[  2] = ['NN', 'VBD', 'DT', 'NN', '.']  [0, 3, 0, 0, 0]  0.583

Highest Scoring Meaning:
meaningIndices = [0, 2, 0, 0, 0]   score = 0.631
Mary         NN      a woman
enjoyed      VBD     (past tense) derive or receive pleasure from   verbBase = enjoy
the          DT      the definite article, used to specifically identify something
```

```
movie        NN      a recording of moving images that tells a story and that people watch on a screen or television
.            .       .

backTrace Records:
Mary -> woman -> _female_ -> accept -> receive -> enjoyed
Mary -> woman -> _human_ -> watch -> movie
enjoyed -> pleasure -> _people_ -> movie
```

Here are a few additional examples:

```
Time flies like an arrow.

Top 3 Meanings out of 310 calculated:
sortedMeaningsAndScores[  0] = ['NN', 'VB', 'IN', 'DT', 'NN', '.']  [0, 3, 0, 0, 0, 0]  1.844
sortedMeaningsAndScores[  1] = ['NN', 'VB', 'IN', 'DT', 'NN', '.']  [0, 2, 0, 0, 0, 0]  1.786
sortedMeaningsAndScores[  2] = ['NN', 'VB', 'IN', 'DT', 'NN', '.']  [0, 5, 0, 0, 0, 0]  1.750

Highest Scoring Meaning:
meaningIndices = [0, 3, 0, 0, 0, 0]   score = 1.844
Time         NN      the continuum of experience in which events pass from the future through the present to the past
flies        VB      to float, wave, or soar in the air
like         IN      having the characteristics of
an           DT      the indefinite article, used to non-specifically identify something, when followed by a word starting with a
vowel
arrow        NN      a projectile with a straight thin shaft and an arrowhead on one end and stabilizing vanes on the other
.            .       .

backTrace Records:
Time -> _continuum_ -> float -> flies
Time -> future -> delivery -> _like_
Time -> events -> point -> _arrow_
flies -> wave -> _like_
flies -> _air_ -> vanes -> arrow
like -> shaft -> arrow


Carry on my wayward son.


Top 3 Meanings out of 300 calculated:
sortedMeaningsAndScores[  0] = ['VB', 'IN', 'JJ', 'JJ', 'NN', '.']  [0, 0, 0, 0, 0, 0]  0.750
sortedMeaningsAndScores[  1] = ['VB', 'IN', 'UH', 'JJ', 'NN', '.']  [0, 0, 1, 0, 0, 0]  0.750
sortedMeaningsAndScores[  2] = ['VB', 'IN', 'JJ', 'JJ', 'NN', '.']  [2, 0, 0, 0, 0, 0]  0.631

Highest Scoring Meaning:
meaningIndices = [0, 0, 0, 0, 0, 0]   score = 0.750
Carry        VB      to support or hold in a certain manner
on           IN      used as a function word to indicate position in contact with and supported by the top surface of
my           JJ      of or relating to me or myself especially as possessor, agent, object of an action, or familiar person
wayward      JJ      resistant to guidance or discipline
son          NN      a male human offspring
.            .       .

backTrace Records:
Carry -> _hold_ -> container -> top -> on
Carry -> hold -> activity -> _human_ -> son
on -> top -> head -> _human_ -> son
```

Note that the above example “Carry on my wayward son” has not correctly classified the word “on”. It has been classified as a preposition, when it should be classified as an adverb, as used in the adverb phrase “carry on”, to continue or persist against adversity. This type of phrase is present in the dictionary, but the current system does not yet take it into account.

4. Discussion

Clearly, the present project is a preliminary Proof of Concept. It is only tested on a small number of example sentences, about 20. It requires a complete dictionary which includes all words that it can encounter in its inputs, and also includes all words used in its own definitions. Currently, the dictionary does not meet those requirements, as shown in Figure 5.

File	Items	Meanings	Nouns	Verbs	Adject	Adverbs	Pronoun	Prep	Article	Conj	Interj	Phrases	First Entry
&.json	21	27	17	5	5	0	0	0	0	0	0	0	'Ã¦sthetic'
a.json	7096	9848	5969	1003	2332	515	0	13	2	2	1	11	'a cappella'
b.json	6794	10198	6811	1593	1545	240	0	4	0	4	0	1	'BC'
c.json	10697	16394	10860	2692	2513	325	0	0	0	0	3	1	'CE'
d.json	5673	9053	5274	1892	1586	299	0	1	0	0	1	0	'dead'
e.json	3805	5666	3404	903	1126	233	0	0	0	0	0	0	'essentially'
f.json	4795	7680	4867	1244	1276	286	0	6	0	0	1	0	'fundamentally'
g.json	3806	5582	3772	972	731	105	0	0	0	0	0	2	'gravely'
h.json	4335	6198	4165	777	1042	206	7	0	0	0	0	1	'horseback'
i.json	3778	5707	2973	725	1510	488	7	3	0	0	1	0	'it'
j.json	763	1217	826	192	158	41	0	0	0	0	0	0	'just'
k.json	885	1260	910	224	116	10	0	0	0	0	0	0	'kindly'
l.json	4031	5942	3931	753	1069	188	0	1	0	0	0	0	'largely'
m.json	5890	8273	5779	950	1332	209	2	0	0	0	1	0	'more or less'
n.json	2539	3469	2167	247	911	143	0	0	0	0	0	1	'never'
o.json	2678	3883	2238	528	845	258	0	13	0	1	0	0	'o.k.'
p.json	8963	13296	8976	1879	2046	395	0	0	0	0	0	0	'plumb'
q.json	395	689	447	99	112	31	0	0	0	0	0	0	'quite'
r.json	4928	8443	5103	1913	1230	197	0	0	0	0	0	0	'roughly'
s.json	12717	20217	12876	3448	3313	579	1	0	0	0	0	0	'scarcely'
t.json	5698	8761	5636	1484	1319	305	7	4	1	4	0	1	'the'
u.json	2882	3990	981	350	2341	309	0	7	0	2	0	0	'unsatisfactorily'
v.json	1440	2192	1424	262	435	71	0	0	0	0	0	0	'vastly'
w.json	3065	4713	3187	673	711	121	5	10	0	5	0	1	'when'
x.json	107	113	57	3	53	0	0	0	0	0	0	0	'xenophobic'
y.json	225	368	238	51	55	16	5	0	0	1	1	1	'yet'
z.json	171	234	161	25	44	4	0	0	0	0	0	0	'zealously'
Total	108177	163413	103049	24887	29756	5574	34	62	3	19	9	20	

```
totalUniqueDefWordsInDictionary      =   24613
totalUniqueDefWordsOutsideDictionary =   11045
totalUniqueDefWords                  =   35658
```

Figure 5: Properties of the current dictionary.

The dictionary was initially based on the open-source WordSet dictionary [14], a 56 MB github repository consisting of 27 json files. It appears that it was originally based on the dictionary part of WordNet [13], which consisted of nouns, verbs, adjectives and adverbs, with no pronouns, prepositions, articles, or conjunctions, and then some additional definitions were added (largely scientific and social terms). In order to be useful for the current model, we had to add pronouns, prepositions, articles and conjunctions, and add and correct many definitions.

The current dictionary has 108,177 entries, with a total of 163,413 meanings. 103,049 of those meanings are nouns, 24,887 are verbs, 29756 are adjectives, and 5,574 are adverbs. The total number of unique words used in the definitions is 24,613. There are 11,045 unique words used in the definitions that do not appear in the dictionary. That's a big problem. In order to have a complete dictionary, we would have to add 11,045 new words and their corresponding definitions. That would either require acquiring or licensing a more thorough commercial dictionary, or otherwise manually adding the remaining 11,045 new entries.

Another resource is needed to make the system work. Recall that the graph theory distance calculation was performed on "nontrivial" word pairs (excluding the word "the" in the given example). In general, we need a list of "trivial" words that should be excluded in the distance calculation, otherwise all the distances will be 2 because matches will be found with trivial words like "and", "or", "the", "of", etc. In the current system, we manually built up a trivial word list as needed to make the examples work.

Note that the current model effectively does an exhaustive search across all possible combinations of word meanings. This is practical for small input sentences. For much longer input sentences, it will be necessary to make many performance optimizations: pre-computing and storing frequently-used distance measures, pruning the search tree, carefully avoiding redundant computations, subdividing the inputs into manageable phrases and operating separately on them, etc. Note that most other approaches such as Stanford CoreNLP use a serial pipeline for tokenizing, parsing, word sense disambiguation, etc. Such approaches are less demanding computationally than exhaustive search, but can fail badly if an error is made in an early processing stage. The current approach can achieve an optimal parsing and word sense disambiguation by jointly solving both problems simultaneously via exhaustive search.

Finally, imagine that we were going to really build a full working system, capable of working on arbitrary real-world examples, properly benchmarked and optimized for best performance. We would need an extensive set of labeled examples to use as a training data set, additional labeled examples to use as a test data set, and an automatic symbolic learning rule to adapt the key resources (dictionary entries, trivial word list, and distance metric definition) to maximize the performance on the training and test data sets. This would be an exciting but challenging problem to work on, potentially for a fairly large team of professional engineers. This could be many engineer-years of effort, definitely beyond the scope of the current Proof of Concept project.

5. Conclusions

We have given a Proof of Concept demonstration of a Computational Model of Natural Language Understanding, using a structured dictionary, a graph theory distance metric, an exhaustive search across all possible sentence meanings, and backtrace of the best paths to provide explainability.

References

[1] Edoardo Barba, Luigi Procopio and Roberto Navigli (2021), "ConSeC:Word Sense Disambiguation as Continuous Sense Comprehension", *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing,* pages 1492–1503.

[2] Alessandro Raganato, Jose Camacho-Collados, and Roberto Navigli (2017). "Word sense disambiguation: A unified evaluation framework and empirical comparison", in *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*: Volume 1, Long Papers, pages 99–110, Valencia, Spain. Association for Computational Linguistics.

[3] Gary Marcus, (March 10, 2022), "Deep Learning Is Hitting A Wall", https://nautil.us/deep-learning-is-hitting-a-wall-14467/

[4] Davis, E. & Marcus, G. GPT-3, "Bloviator: OpenAI's language generator has no idea what it's talking about." MIT Technology Review (2020).

[5] Walid Saba, (March 1, 2022), *Sorcero Employee Spotlight: Dr. Walid Saba.* https://www.sorcero.com/blog/employee-spotlight-walid-saba

[6] Lloyd Watts, (2012) "Reverse-Engineering the Human Auditory Pathway", in *J. Liu et al. (Eds.): IEEE WCCI 2012, LNCS 7311*, pp. 47–59, 2012.

[7] William Matchin and Gregory Hickok, (March 2020), "The Cortical Organization of Syntax", *Cerebral Cortex,* March 2020; 30: 1481–1498.

[8] Binder JR, Desai RH, (2011). "The neurobiology of semantic memory", *Trends Cogn Sci.* 15:527–536.

[9] Lloyd Watts, (2022), The Architecture of the Human Brain. http://lloydwatts.com/images/wholeBrain_007.jpg

[10] Michael Lesk (1986), "Automatic Sense Disambiguation Using Machine Readable Dictionaries: How to Tell a Pine Cone from an Ice Cream Cone", In *SIGDOC '86: Proceedings of the 5th annual international conference on Systems documentation*, pages 24-26, New York, NY, USA. ACM.

[11] Roberto Navigli. "Word Sense Disambiguation: A Survey", *ACM Computing Surveys*, 41(2), 2009, pp. 1–69.

[12] Bidirectional Search Wikipedia page. https://en.wikipedia.org/wiki/Bidirectional_search

[13] WordNet Wikipedia page. https://en.wikipedia.org/wiki/WordNet

[14] WordSet dictionary. https://github.com/wordset/wordset-dictionary