Hallucinations in Garden Path sentences

Alankrita Bhatt[1], Sharbatanu Chatterjee[2]
[1] *Department of Electrical Engineering*
[2] *Department of Computer Science and Engineering*
Indian Institute of Technology Kanpur

November 21, 2014

Abstract

Garden path sentences are those in which an ambiguity in the structure biases the reader's interpretation so strongly toward a wrong interpretation that when a disambiguating word is encountered, the meaning is very tough to recover, if at all possible. In this project, we study sentences that are not traditionally garden path, but that are interpreted as garden path because the initial input, i.e. the initial part of the sentence is perceived in a "distorted" manner, since this erroneous reporesentation correlates well with the part of the sentence immidiately after. Hence, we prove that the linguistic system employs a feedback system from the current input to the input before, and actively revises the hitherto perceived input. We also show that linguistic representations made upon re-analysis of ambiguous sentences are often only partially correct. We present two behavioural experiments that confirm our hypotheses.

1 Introduction

Garden-Path sentences are sentences, quite simply put, which lure the reader "down the garden path", that is those which lure the reader into parsing the sentence in a way that turns out to be a cul-de-sac. For example, consider the most famous garden-path sentence:

The horse raced past the barn fell.

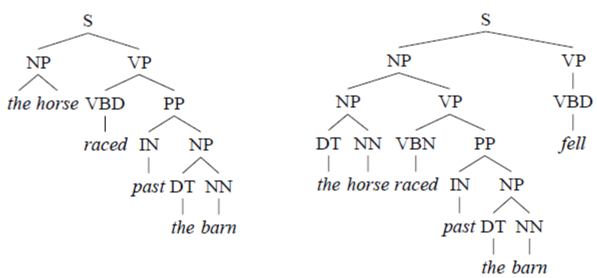

Figure 1: Most likely interpretation (L) vs. Correct interpretation(R)

The reader will be very strongly biased towards interpreting this sentence as one which has *raced* as the MV (Main Verb) as opposed to one where it is the beginning of a RR (Reduced Relative) clause. Subsequently on encountering the word *fell*, the reader will have to go back and analyse the sentence again.
In this project, we follow the work done in [1], and show that in sentences where the perceptual neighbour of the initial part has strong correlation with the rest of the sentence, the cognitive system percieves the perceptual neighbour of the previous part of the sentence and not the true raw input. These "hallucinations" then can lead to signs exhibited by traditional garden-path sentences at a disambiguating word. Hence, we show that linguistic information is used both proactively and retroactively. We further show that the reader, upon encountering this disambiguating word, does not reconstruct the entire parse tree and settles for a somewhat correct parse tree, thus validating the "good-enough" theory of [4].
We also use the surprisal theory of [2], [3] to show that the cognitive effort, measured by the time to read a particular word is linearly related to the surprisal.
Section 2 reviews some theories of statistical language and language processing.
Section 3 explains the behavioural experiments used to validate our hypotheses.
Section 4 shows the obtained results and makes inferences. Section 5 concludes.

2 Background

2.1 Statistical Language

The cognitive effort can be quantized in terms of an information theoritic measure called the "surprisal" or the "Shannon information content", defined as The surprisal theory suggests that the cognitive effort in reading a word w_i of a sentence is defined by

$$Effort(w_i) = \log \frac{1}{P(w_i|w_{1,2...i-1}, C_{txt})} \tag{1}$$

Here, C_{txt} refers to the context of the sentence. In the psycholinguistic literature, it has been proved that the reading time for a particular word (which may be taken as the cognitive effort) and the suprisal are linearly related([6]). The surprisal theory has achieved considerable success because it matches well with emperical results.

Here, the probablities are obtained by using the PCFG (Probabilistic Context Free Grammar) probabilies from a parsed corpus.
The PCFG probabilities can also be used to compute the probability of a sentence. The probability of a sentence is simply the product of all the rules used to generate the sentences.

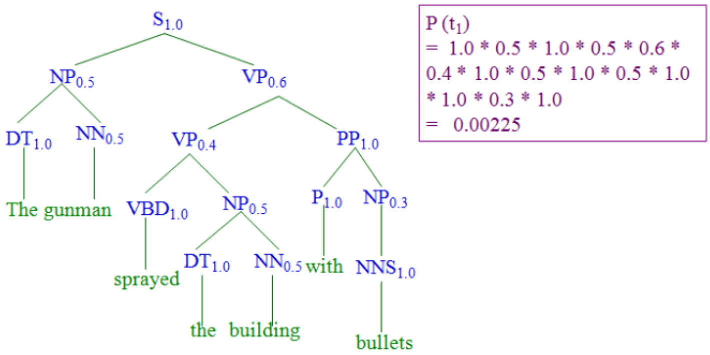

Figure 2: Computing the probability of a sentence using assumed PCFG probabilities[5]

TOP	→ S .	1.000000
S	→ INVERTED NP	0.003257
S	→ SBAR S	0.012289
S	→ SBAR , S	0.041753
S	→ NP VP	0.942701
INVERTED	→ PP VBD	1.000000
SBAR	→ INSBAR S	1.000000
VP	→ VBD RB	0.002149
VP	→ VBD PP	0.202024
VP	→ VBD NP	0.393660
VP	→ VBD PP PP	0.028029
VP	→ VBD RP	0.005731
VP	→ VBD	0.222441
VP	→ VBD JJ	0.145966
PP	→ IN NP	1.000000
NP	→ DT NN	0.274566
NP	→ NNS	0.047505
NP	→ NNP	0.101198
NP	→ DT NNS	0.045082
NP	→ PRP	0.412192
NP	→ NN	0.119456

Figure 3: Example probabilities obtained from the parsed Brown corpus

2.2 Previous Theories of Sentence Processing

Incremental Processing Theory: The incremental processing theory states that upon enountering a word, the processor immidiately settles upon a semantic and syntactic interpretation. When there is some further input disconfirming the current interpretation, the processor has to go back and analyze the sentence again. The opposite of an incremental processor would be some sort of a top-down processor, which creates a bottom up parse tree, instead of sequentially building one as the sentence progresses. As per this theory, the current input is used only to predict the future input.

Good-Enough theory: According to the good-enough theory of [4], the linguistic representations of a sentences are not always complete and accurate, and often people settle for representations that are not necessarily reflective of the true content of a sentence. Language representations are often partial and semantic analysis is often incomplete. So, upon encountering an ambiguity, complete reanalysis does not happen and the cognitive system settles for a parse tree that is just "good enough".
We aim to show that the good-enough theory holds and also to show that the incremental-processing theory needs to account for a feedback mechanism from the current input to the already read input.

3 Experimental Setup

We performed 2 behavioural experiments to confirm our hypotheses. Our subjects for the reading times experiment were 35 undergraduates, male and female, from IIT Kanpur. Our subjects for the Gaze tracking experiment were 34 people from the SE367 course.

3.1 Reading Times

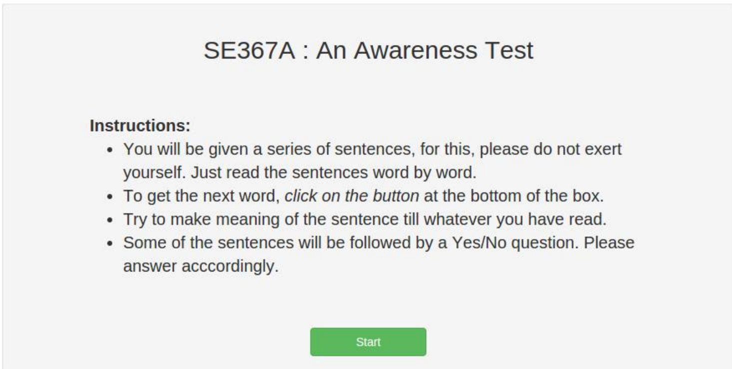

Figure 4: Instructions for the reading times experiment

The above-shown application (http://home.iitk.ac.in/∼sharbatc/se367/cogapp) was developed and used in the first experiment. The participants had to press a button to reveal the next word in a sentence as shown:

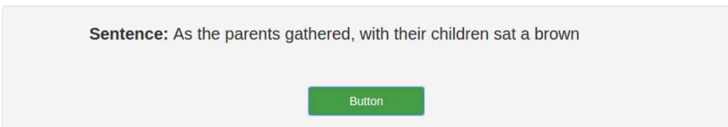

Each sentence was followed by a yes/no sentence to test the semantic comprehension of the sentence by the reader, like

There were 6 experimental items. The experimental items were interspersed with filler sentences, thus giving a total of 16 sentences to read.
The time between two subsequent button presses was recorded. This was taken to be a measure of the time taken to read that particular word.

3.2 Gaze Tracking

We supplemented our reading times experiment with a gaze tracking experiment, along the lines of that described in [7]The instruments used were developed by SensoMotoric Instruments Ltd.
The participants were instructed to pay attention and read. They were told not to memorize it or read too thoroughly, but only to skim through it attentively. We did this because we wanted results in a naturalistic setting. Their saccadic eye movements were recorded.

4 Results and Inferences

We plotted the reading time for each word in the experimental sentences. We also made qualitative inferences from the gaze tracking data.

4.1 Reading Times

The following are some of the observed results for the experimental sentences:

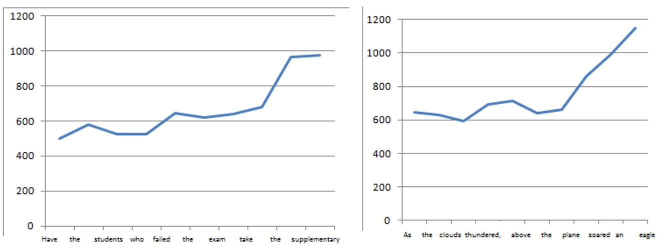

Figure 5: Classic (L) vs. Hallucinated (R) garden path disambiguation

Note the similarities in human reactions to classic and hallucinated garden path sentences. This effectively proves our theory of retroactive interference - had the input comma been retained in the cognitive system, this wouldn't have happened, as *soared* can be easily accomodated in the second clause in the sentence.

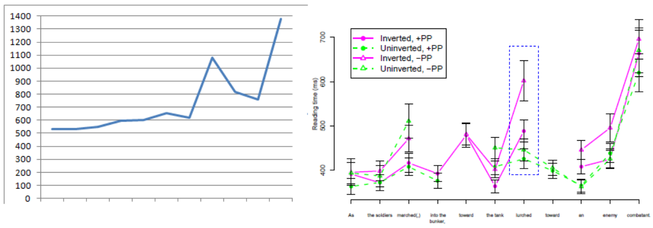

Figure 6: Our result (L) vs. Levy's result (R)

The answers to the yes/no comprehension questions posed also helped us gain insight into the interpretation of the sentence by the reader.

Table 1: Answers to the yes/no questions

Sentence	Question	Yes(%)	No(%)
She swept the flour.	Was the floor swept?	54	46
As the clouds thundered,above the plane soared an eagle.	Did the clouds thunder above the plane?	36	64
Lose the knot that was made.	Are you instructed to loosen a knot?	71	29
As the soldiers marched, towards the tank lurched an enemy combatant.	Did the soldiers march toward the tank?	26	74
We painted the wall with cracks.	Did we paint the wall with colour?	31	69
I told her children are noisy.	Are her children noisy?	68	32

The first question shows us that linguistic information is used proactively, which is known. The third sentence and the strong percentage of error (71%) shows us that the information is also used retroactively, as the cognitive system, owing to the high correlation between *loose* and *knot* corrects *lose* to *loose*. Sentences 2 and 4 are hallucinated garden paths. Sentences 5 and 6 show that humans do indeed employ a "good-enough" theory and don't question the possibility of painting a wall (using) cracks.

The parsed tree of the hallucinated garden path sentence is shown below (Parsed using the Berkley parser)

```
(ROOT
  (FRAG
    (SBAR (IN As)
      (S
        (NP (DT the) (NNS soldiers))
        (VP (VBD marched))))
    (, ,)
    (SBAR (IN towards)
      (S
        (NP (DT the) (NN bunker))
        (VP (VBD lurched)
          (NP (DT an) (NN enemy) (NN combatant)))))
    (. .)
    (NP (NNP *CR*))))
```

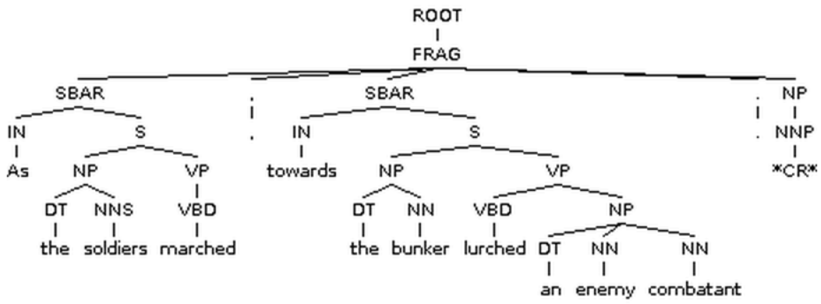

Figure 7: Parse Tree

4.2 Results from the gaze tracking experiment

We first present a frame by frame view of the results of our gaze tracking experiment, on the sentence *As the soldiers marched, towards the bunker lurched an injured enemy combatant.* The radius of the circle around each word is proportional to the saccade time.

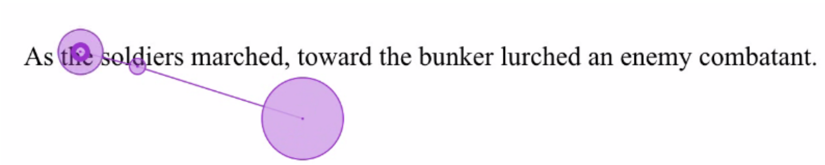

Figure 8: Saccade 1

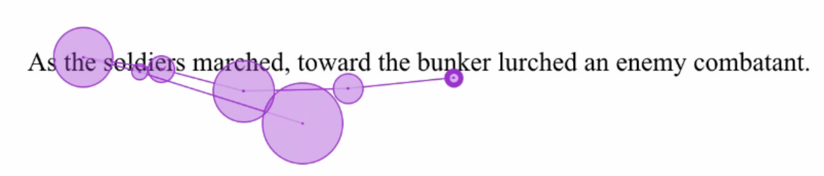

Figure 9: Saccade 2

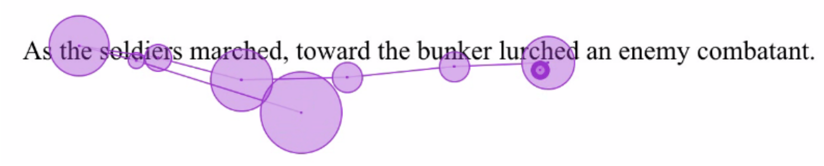

Figure 10: Saccade 3

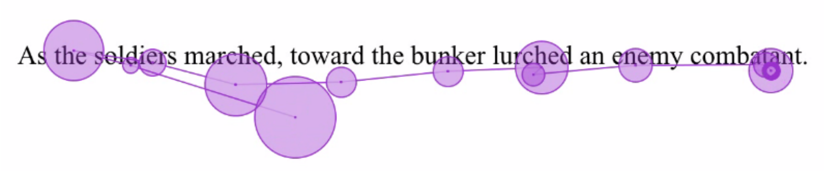

Figure 11: Saccade 4

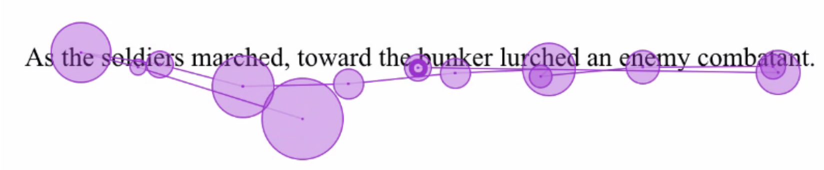

Figure 12: Saccade 5

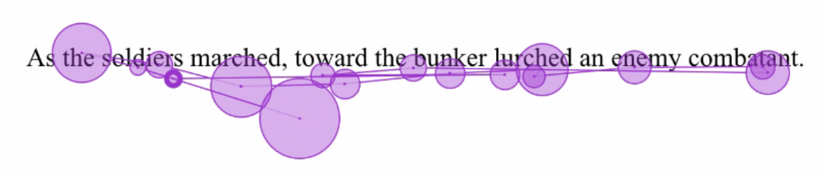

Figure 13: Saccade 6

Qualitatively, using the data from the saccades, we can figure out from the eye movement about the mental processes that might be going on in the reader as he/she goes through the sentence. The given example is one of many which have been used from the pool of data that was collected from the gaze tracking experiment. The large purple circle representing the initial fixation of the eyes that is ubiquitous in every figure is to be ignored. The rest of the data shows how the gaze travels across this typically hallucinating garden path sentence as the person reads through. Saccades 1 through 4 show a linear traversal throughout the sentence while the later figures show how the eye goes back to the hallucinating area (Saccades 5 and 6) where the ambiguation occurs. It is noticed that the disambiguating comma is skimmed over, even on reevaluation of the sentence, providing a basis for our assumption that the garden path effect can be hallucinated.

5 Conclusions and possible future work

5.1 Conclusions

We conclude that the human cognitive system is indeed optimized to use the linguistic information optimally. It is also "lazy" so to say, and does not construct parse trees that are completely correct, instead being satisfied with partially correct "good enough" representations.

5.2 Future Work

Garden-path sentences in Indian languages have not been studied. However, it is tough to find such sentences. This may possibly be because of the SOV (Subject-Object-Verb) structure of most Indian languages, which makes it tough to induce a Noun/Verb ambiguity. Some of the examples we were able to find were of a poetic nature, which allowed the writer to slightly change the structure of the sentence.

এসেছি যা ফেলে তার চিহ্ন তবু ছাড়ে না অভ্যাস
সমস্ত চিতকারশব্দে শুনতে পাই স্তব্ধতার শ্বাস।

কাদম্বিনী মরিয়া প্রমাণ করিল যে সে মরে নাই।

Figure 14: Bengali Garden Path Sentences

It would be interesting to study Garden-pathing in other languages, since it would help in answering questions like:

- Do speakers of language L1 parse L1 the same way as speakers of another language L2 do?
- Will a native speaker of L2, upon reading a garden-path sentence in L1, employ the same amount of cognitive effort as a speaker of L1? This may not be true, considering what is an unusual or low-probability structure in L1 may not be a low-probability sentence structure in L2.

An interesting case study in this regard would be the study of the Romani language, an Indo-European language

6 Acknowlegements

We thank Prof. Amitabha Mukerjee for exposing us to the interesting possibilities and questions in the field of Cognitive Science, and for his invaluable support. We also thank Rahul Meena for his help with the gaze-tracking system, without which the project would not have been complete. We also thank Saurav Kumar, for his help in developing the application for the reading times experiment.

7 References

[1] Levy, Roger. "Integrating surprisal and uncertain-input models in online sentence comprehension: formal techniques and empirical results." Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1. Association for Computational Linguistics, 2011.

[2] Hale, John. "A probabilistic Earley parser as a psycholinguistic model." Proceedings of the second meeting of the North American Chapter of the Association for Computational Linguistics on Language technologies. Association for Computational Linguistics, 2001.

[3] Levy, Roger. "Expectation-based syntactic comprehension." Cognition 106.3 (2008): 1126-1177.

[4] Ferreira, Fernanda, Karl GD Bailey, and Vittoria Ferraro. "Good-enough representations in language comprehension." Current Directions in Psychological Science 11.1 (2002): 11-15.

[5] Speech and Natural Language Processing and the Web/Topics in Artificial Intelligence Programming, Pushpak Bhattacharya - Lecture 16 (http://www.cse.iitb.ac.in/ cs626-449/slides/Lec%2016.pdf)

[6] Smith, Nathaniel J., and Roger Levy. "Optimal processing times in reading: a formal model and empirical investigation." Proceedings of the 30th annual conference of the cognitive science society. 2008.

[7] Slattery, Timothy J., et al. "Lingering misinterpretations of garden path sentences arise from competing syntactic representations." Journal of Memory and Language 69.2 (2013): 104-120.