

A Comparative Study on Ambiguity Resolution in Garden-Path Sentences Between Humans and Artificial Intelligence

Jiahao Liu

Beijing International Studies University, Beijing, China

Abstract: The garden-path sentence is a paradigm for studying cognitive mechanisms. This study has found that both humans and artificial intelligence encounter difficulties in garden-path sentences. Humans face these difficulties due to limited working memory, while artificial intelligence relies on probabilistic choices based on big data. This research provides a theoretical foundation for further investigation into humans' cognitive processing patterns and for the improvement of the next generation of artificial intelligence to make its cognitive models true-to-human, and for a presentation of a novel comparative research paradigm for studying human and artificial intelligence.

Keywords: Syntactic Ambiguity Resolution; Garden-Path Sentence; Comparison Between Human and Artificial Intelligence; Large Language Model



Copyright © 2025 by author (s) and SciScan Publishing Limited

This article is licensed under a [Creative Commons Attribution-NonCommercial 4.0 International License](https://creativecommons.org/licenses/by-nc/4.0/). <https://creativecommons.org/licenses/by-nc/4.0/>

1 Introduction

With the breakthrough achievements of artificial intelligence (AI), especially large language models (LLMs) based on deep learning models (DLM, e.g., Transformer), in natural language processing tasks, an unprecedented opportunity and challenge have emerged: do these artificial intelligence systems, capable of generating text fluently, truly understand the inherent structure of language? When faced with garden-path sentences that confuse humans, will artificial intelligence imitate humans' error patterns or encounter a completely different computational dilemma? Considering artificial intelligence models as participants provides researchers with a new perspective to re-examine, test, and inspire human cognitive theory (Marvin, & Linzen, 2018).

This study adopts the research paradigms of Computational Cognitive Linguistics, Computational Psycholinguistics, and Computational Neurolinguistics, which not only utilise artificial intelligence to predict human behaviour but also study the two types of intelligent agents in the same problem domain, thereby enabling bidirectional mapping and interpretation between computational principles and biological mechanisms. This provides a methodological paradigm for future comparative studies on language cognition on a broader scale.

2 Cognitive dilemmas of humans: prediction, conflict, and re-evaluation with limited resources

For humans, the processing dilemma of garden-path sentences stems from the inherent limitations of their cognitive architecture. Theories such as Serial Process and Constraint Satisfaction suggest that this is a process where sequential prediction failure triggers cognitive conflict and requires significant resources for reanalysis. Humans' language comprehension is highly predictive, and readers make instant predictions based on syntactic cues and lexical preferences. Once subsequent input violates this prediction, the garden-path effect occurs. This conflict manifests at the neural level as specific electroencephalographic components, such as P600.¹ Recent research has further linked this process to cognitive control brain networks (especially the left inferior frontal gyrus), emphasising its central role in inhibiting erroneous interpretations and resolving ambiguity conflicts. Successful reanalysis heavily relies on limited cognitive resources such as working memory (context window utilisation) and executive function. Differences in individual working memory capacity directly affect the efficiency of recovery from the dilemma of the garden-path sentence. Therefore, the human dilemma is a conscious correction process that requires effort due to conflicts between active prediction and subsequent input in a cognitive system with limited sequence processing and resources. The cognitive processing of syntactic ambiguity resolution for a garden-path sentence in humans can be summarised into the following three stages:

2.1 Quick prediction based on clues

When readers begin to process a garden-path sentence (e.g., "The horse raced past the barn fell."),² They do not wait until they have read the entire sentence before analysing it; instead, they process it as they read. When they come across "the horse raced," their brain, based on lexical preference and the Minimal Attachment Principle, treats "raced" as the main verb.³ This is much simpler than analysing it as a relative clause with an omitted relative pronoun ("the horse that was ridden past the barn"). As a result, a mental model of "this horse racing past the barn" is quickly established.

2.2 Conflict monitoring and processing interruption

When readers encounter that critical disambiguating word, they are faced with a dilemma that arises when they read "fell." This new verb cannot be integrated into the existing subject-predicate structure. The sentence lacks a subject ("the horse" has already been occupied), making it grammatically incoherent. The conflict monitoring system in the brain (mainly related to the prefrontal cortex, especially the anterior cingulate cortex) is strongly activated. It detects a serious discrepancy between the current input and the initial predictive model. This conflict will immediately manifest as a significant increase in reading time—specifically, the reader's gaze lingers longer on the word "fell." The reader will trace back to the first half of the sentence and re-examine the word "raced." In event-related potential (ERP) studies, a

¹This is a positive wave that appears ca. 600 ms after the stimulus, and is widely recognised as a clear electrophysiological marker of syntactic integration difficulties and syntactic reanalysis.

²This is a classic example sentence proposed by Bever (1970).

³Based on lexical preference, since "horse" is a typical subject and "raced" is a high-frequency verb, the brain will prioritise parsing it as a "subject-predicate" structure, i.e., this horse races. The Minimal Attachment is a classic principle in psycholinguistics, which states that the brain tends to construct the simplest structure with the fewest required syntactic nodes (Frazier & Fodor, 1978).

typical brain electrical component, P600, will be observed.

2.3 Cognitive control and reanalysis

After conflict monitoring, this is the stage that consumes the most cognitive resources and also the key stage of resolving dilemmas. The cognitive control network of the brain needs to suppress that powerful but erroneous initial interpretation (e.g., “This horse is racing.”). This process requires the involvement of executive function. Working memory is fully mobilised, and readers must maintain all the words (e.g., the, horse, raced, past, the, barn) that have appeared in memory. Rearrange “raced” from its primary verb interpretation to its past participle form. Reanalyse the sentence structure from [The horse] [raced past the barn] to [The horse [raced past the barn]] [fell] (i.e., The horse that was ridden past the barn fell.) (as shown in Figure 1) (Hale, 2014). This “dismantling and reconstruction” process heavily relies on working memory and executive control abilities. Research has shown that individuals with lower working memory capacity have more difficulty completing the reanalysis, require a longer time, and have a higher failure rate (Just, & Carpenter, 1992).¹




Figure 1 Tree diagram of comparison of garden-path structure and correct structure

The dilemma faced by human in garden-path sentences stems from the contradiction between their efficient yet arbitrary cognitive strategies and limited cognitive resources. The reason for encountering this dilemma is that the brain, as a powerful prediction machine, sacrifices absolute caution for speed. The reason for getting into this dilemma is that recovering from errors requires the use of limited working memory and cognitive control resources to execute a time-consuming reanalysis process. Therefore, the garden-path sentence is like a cognitive prism, which clearly magnifies and presents those cognitive processes (e.g., prediction, conflict, inhibition, reconstruction, etc.) in human language processing that are usually rapid, automated, and imperceptible. This is fundamentally different from the probabilistic errors of artificial intelligence, which are based on global statistics and lack true cognitive conflict.

3 The simulation dilemma of artificial intelligence: architecture differences and “pseudo-understanding” driven by statistics

Currently, although research has begun to test the performance of artificial intelligence in processing garden-path sentences, most works remain at the simple comparison of behavioural performance, determining whether it is right

¹This is a highly mature and widely validated hypothesis in psycholinguistics and cognitive psychology, originating from the Capacity Theory proposed by Just and Carpenter (1992). It posits the existence of a general speech working memory system with limited capacity, which directly constrains language processing.

or wrong, lacking a systematic comparison framework that delves into the cognitive processing mechanism. There is relatively little research on the fundamental similarities and differences in attention allocation, utilisation of working memory, and strategies for resolving syntactic ambiguity between humans and artificial intelligence when facing the same linguistic dilemma. Therefore, this article will delve into these issues.

Through the counterfactual model of artificial intelligence, humans' cognitive theories can be tested. If artificial intelligence replicates humans' behaviour through pure statistical learning, it challenges the innate hypothesis of Chomsky's Formal Languages that humans are born with universal knowledge of language and construct a symbolic derivation system mathematically. If its error patterns differ significantly from those of humans, it may highlight the central role of human sequence processing and embodied experience based on limited resources (Bisk, Holtzman, & Thomason, et al., 2020).

A deep analysis of the failure cases of artificial intelligence in garden-path sentence can accurately diagnose the cognitive defects of current artificial intelligence models (e.g., the lack of true reanalysis ability, symbol grounding problems, etc.), thus providing theoretical guidance for building a more cognitively reasonable and robust next-generation artificial intelligence model (e.g., introducing an explicit working memory module or central control mechanism).

For artificial intelligence, especially LLMs based on deep learning models (e.g., Transformer), the root of their dilemma lies not in insufficient resources but in their data-driven statistical nature and parallel computing architecture characteristics, which lead to a distinct ambiguity resolution pattern from humans. The self-attention mechanism of Transformer allows it to consider all information in the entire sentence when processing any word, which contrasts with humans' sequence processing patterns. Theoretically, this should make artificial intelligence less susceptible to being misled by syntactic ambiguities. However, research has shown that the errors of artificial intelligence stem from its excessive reliance on the global statistical patterns of training data. If a certain ambiguous structure dominates statistically in the data, the model will still probabilistically choose the wrong path (Hu, & Levy, 2023). Artificial intelligence does not reanalyse in the same way as humans do. When the input sequence changes from "The horse raced past the barn" to a complete sentence containing "fell," the model does not backtrack and correct the initial structure, but performs a forward computation (forward computation refers to the process where input data are passed through the network layer by layer, ultimately resulting in output results) based on the entirely new input sequence. Its output result is an emergent property after parallel computation of all information, rather than a conscious corrective action (Pandia & Ettinger, 2021).

The following are the three stages of the processing mechanism and dilemmas arising from artificial intelligence:

3.1 Training rather than understanding—the formation of statistical bias

During the training phase, artificial intelligence models (e.g., the GPT series) learn vast amounts of text data. However, they do not learn grammatical rules; by contrast, they learn the co-occurrence probabilities between words. When the model identifies "the horse raced" among billions of sentences, it discovers that the probability of "raced" being the main verb immediately following "horse" is much higher than that as a past participle. This high-frequency pattern forms a profound statistical bias in the model's parameters. Artificial intelligence does not have the concept of grammatical categories such as subject and predicate; it only has contextual collocation possibilities embodied through vectors (i.e., symbols) and attention weights.

3.2 Calculation rather than prediction—probabilistic preference in parallel processing

When artificial intelligence receives a garden-path sentence, its process is fundamentally different from humans' sequential processing. Thanks to the self-attention mechanism, Transformer-based artificial intelligence models consider all words in the entire sentence when processing any single word. Theoretically, it can identify the “fell” at the end of the sentence from the very beginning. However, despite being able to identify all the information, the core of artificial intelligence's decision-making is based on the entire input to generate the next word with the highest probability (i.e., calculating the probability of the entire sequence). When processing “the horse raced past the barn.” However, “fell” is next to it, owing to the strong statistical inertia within the model, it may still assign a very high probability to the subject-predicate structure. Its dilemma lies in the fact that global information may not necessarily overwhelm the strong local statistical preferences formed during training.

3.3 Coverage rather than reanalysis—the essence of forward computation

This is the fundamental difference from humans. The artificial intelligence does not have moments of “insight” or “retrospection.” Artificial intelligence's reasoning is a forward computation process. It does not maintain a temporary syntactic structure in its working memory as humans do. For the model, the input sequences “The horse raced past the barn.” and “The horse raced past the barn fell.” are two different inputs. When the input contains “fell,” the model does not revise its previous understanding of “raced.” Instead, it conducts a global computation again based on this new, longer input sequence. At this point, due to the addition of “fell,” the statistical characteristics of the entire sequence have changed, and the model may output the correct interpretation. However, this is not a reanalysis, but a new and different computation that overwrites the previous one. It does not realise that there was something wrong, but simply processes the new input.

Although LLMs can achieve or even surpass humans' performance on many tasks, their performance is weakly correlated with humans' in tasks requiring systematic syntactic reasoning (Misra, 2022). This implies that the models may solve problems through shallow heuristic strategies rather than deep syntactic comprehension. When faced with structures as garden-path sentences, which necessitate deep syntactic analysis, their flawed pseudo-understanding becomes evident. Therefore, the dilemma of artificial intelligence lies in a statistical model lacking true intention and understanding, processing language with its inherent parallel architecture, exhibiting probabilistic biases and implicit computational coverage. Its errors are the result of computation, rather than cognitive conflicts.

Artificial intelligence (especially LLMs) presents a dilemma in disambiguating garden-path sentences, whose façade is similar to humans' cognition but actually has a fundamentally different origin. The core of this dilemma lies in the fact that artificial intelligence's processing is based on forward computation of statistical probabilities, rather than psychological activities involving comprehension and cognitive resources. The dilemma faced by artificial intelligence does not stem from human-like comprehension errors but rather results from its data-driven nature and parallel processing architecture.

4 Connecting two dilemmas: prediction, neural correlation, and bidirectional theoretical implications

When processing garden-path sentences, both human and artificial intelligence seemingly encounter dilemmas,

but the underlying mechanisms that lead to these dilemmas exhibit fundamental similarities and differences. Both will initially misunderstand garden-path sentences. This indicates that humans' experience-based predictions and artificial intelligence's statistics-based pattern matching are common strategies for efficient language processing, but they are also common sources of errors. The dilemma faced by humans stems from their limited rational, biological, and cognitive architecture, while the dilemma faced by artificial intelligence stems from its computationally statistical architecture, lacking understanding. The mechanisms behind these errors are totally different.

The similarity between human and artificial intelligence in resolving ambiguity in garden-path sentences lies in the fact that both encounter difficulties with such sentences and cannot immediately provide the correct interpretation. Both human and artificial intelligence rely heavily on prior experience to guide their parsing. Human relies on life experience and lexical preferences formed through language acquisition, while artificial intelligence relies on statistical frequencies present in training data. After receiving complete information, both human and artificial intelligence can usually arrive at the correct meaning of the sentence. Human does so through effortful reanalysis, while artificial intelligence does so through recalculation based on the complete sentence. The essential differences between the two types of difficulties are specifically reflected in the following five aspects (summarised in Table 1):

4.1 Processing mechanism

Humans engage in sequential processing, whereas artificial intelligence employs parallel processing. When humans process a sentence, their eyes move from word to word, and their brains incrementally construct syntactic structures, which makes them prone to being misled by local information. Artificial intelligence (such as a Transformer) employs parallel processing. Through the self-attention mechanism, it can recognise all words in the sentence when processing each word. Theoretically, it should not be trapped by local ambiguity, thus its errors stem from a deeper statistical nature.

4.2 Source of difficulties

The dilemma of humans stems from cognitive conflict, while the dilemma of artificial intelligence arises from statistical bias. The dilemma of humans originates from cognitive conflict. When subsequent information (like "fell") significantly deviates from initial predictions, the brain's conflict monitoring system (e.g., the anterior cingulate cortex) will trigger an alert, confusing. The dilemma of artificial intelligence arises from statistical bias. Its error is the most likely continuum calculated based on billions of texts in a corpus. If a certain error structure is statistically dominant in the training data, the model will tend to select it probabilistically. There is no conflict but computation.

4.3 Working memory and utilisation of context window

Humans have dynamic working memory, while artificial intelligence models have static working memory. Humans possess active working memory with limited capacity. During reanalysis, it is necessary to suppress old structures and construct new ones, which is a process that requires effort and consumes cognitive resources. On the other hand, artificial intelligence has a passive context window of fixed length. All words are encoded as vectors and mixed. It does not reanalyse but rewrites. When the input sequence changes, it only performs a brand-new forward calculation.

4.4 Error's nature and recovery mechanism

Humans' errors are temporary, conscious failures of analysis, while artificial intelligence's errors are probabilistic

output deviations. Humans' errors are temporary, conscious failures of analysis. Recovery is a centralised, executive process involving cognitive control (inhibition, transformation, and reconstruction). Artificial intelligence's errors, on the other hand, are probabilistic output deviations. Recovery is a decentralised implicit computational override. The model does not revise its comprehension, but outputs computational results for new inputs.

4.5 Symbol grounding and understanding

Human understanding is embodied, while artificial intelligence faces the symbol grounding problem. Humans' understanding is embodied. Words, e.g., "horse," "running," and "falling," are closely linked to humans' sensory-motor experiences, and when processing sentences, relevant scenarios are simulated in the brain. By contrast, artificial intelligence faces the symbol grounding problem. To it, words are merely symbols (i.e., vectors) without real reference and meaning. Its understanding is an empty symbolic calculus, which explains why it may have grammatically correct sentences but absurd common sense.

Table 1 Comparison of dilemmas between humans and artificial intelligence

| Dimension | Dilemma of human | Dilemma of artificial intelligence |
|---------------------|--|--|
| Causes | Limited sequence | Probability handling lacking understanding |
| Mechanism | Sequence prediction, cognitive conflict, and conscious reanalysis | Global attention, statistical preference, and forward computation coverage |
| Essence of error | Temporary, requiring effort to overcome parsing failure | Probabilistic output bias based on the distribution of training data |
| Process of recovery | Consumption of the suppression and reconstruction of cognitive resources | The cost-free and implicit coverage of new computations |

The performance of artificial intelligence on garden-path sentences strongly reveals its fundamental differences from humans: artificial intelligence deals with statistical associations of symbols (i.e., words). However, these symbols are not connected to real-world experiences and references. It knows that the co-occurrence probability of "horse raced" as a subject-predicate structure is extremely high, but it does not understand what it means for a horse to run or fall. Artificial intelligence does not have a central executive system similar to the prefrontal cortex to actively suppress erroneous interpretations and allocate cognitive resources. Its behaviour is completely determined by forward computation. The sequential nature and limited resources of humans may seem like limitations, but they may be the key to forming a strong generalisation and creativity; the parallel nature and unlimited resources of artificial intelligence may seem like advantages, but they may prevent it from engaging in human-like deep reasoning. By comparing these two dilemmas, the uniqueness of human cognition is more apparent.

Therefore, studying the disambiguation of garden-path sentences can clearly contrast the profound differences between human and artificial intelligence in terms of the essence, architecture, and mechanism of intelligence. Research has found that the degree to which words surprise people, calculated from LLMs, exhibits a high correlation with humans' reading time (Merkx, & Frank, 2021). This suggests that prediction error may be a common core quantitative indicator of humans' cognitive load and artificial intelligence computational uncertainty.¹ Furthermore, some studies have attempted to use the internal representations of LLMs (multimodal models such as CLIP, vector space models, or contextual models such as BERT or GPT) to predict human's brain's activity during language processing (e.g., fMRI

¹ The term cognitive load was proposed by John Sweller in 1988.

signals, ECoG, Broca's area, anterior temporal lobe, etc.) (Tuckute, Sathe, & Srikant, 2024; Caucheteux, & King, 2022; Li, O'Sullivan, & Mattingley, 2022; Taoudi-Benchekroun, Christiaens, & Grigorescu, 2022; Schrimpf, Blank, & Tuckute, et al., 2021; Jain, & Huth, 2018; Wehbe, Murphy, & Talukdar, et al., 2014).¹ This indicates that certain levels of artificial intelligence models may have learnt representations similar to those of brain language processing regions, providing a new approach for validating computational models at the neural level.

For cognitive science, artificial intelligence serves as a counterfactual model, and its failures (e.g., its inability to robustly handle garden-path sentences) strongly support the central role of human embodied cognition, innate learning biases, and working memory in language processing (Baggio, 2022). For artificial intelligence, a deep understanding of human cognitive mechanisms points the way to breaking through the current limitations of artificial intelligence. For instance, introducing explicit and structured memory modules and central control systems into artificial intelligence may enable it to possess more human-like and robust reanalysis capabilities (Linzen, 2020).

5 Conclusion

The dilemma faced by humans in garden-path sentences lies in the limited sequence processing, stemming from inherent cognitive resource limitations; whereas the dilemma of artificial intelligence stems from the infinite statistical processing, resulting from the gap between its architecture and true understanding. For future research, one approach could be to expand across languages and examine the universality and specificity of the dilemmas faced by humans and artificial intelligence models under different syntactic structures. Another approach is to conduct detailed model exploration. With the rapid iteration of artificial intelligence models, it is necessary to perform diagnostic assessments on more advanced architectures (such as models incorporating recursion or explicit working memory) to examine their progress in addressing such cognitive dilemmas. A third approach is to conduct empirical tests. In the future, sophisticated cross-experimental designs could be devised, simultaneously collecting behavioural and neural data such as eye movement, ERP, and fMRI from humans, as well as computational data such as internal activation and attention weights from artificial intelligence models, using the same set of garden-path sentence materials for more direct correlation and comparison. A fourth approach could be to shift from static word lists to fully generative model analysis. A fifth approach is to design more diagnostic experiments to isolate the contributions of syntactic, semantic, and pragmatic information in the disambiguation of human and artificial intelligence. A sixth approach is to utilise technologies such as magnetoencephalography (MEG) and electroencephalogram (EEG) with high temporal resolution to explore the millisecond-level dynamic processes of language processing. A seventh approach is to utilise encoding models to understand the abnormal neural activity in clinical populations, e.g., patients with aphasia and Alzheimer's disease. By continuously deepening this comparative research, not only can researchers gain a deeper understanding of the mysteries of the human mind, but they can also point the way forward for creating truly intelligent machines.

References

- [1] Baggio, G. (2022). *Meaning in the brain*. The MIT Press.

¹ The anterior temporal lobe is a region for multimodal integration.

- [2] Bever, T. (1970). *The cognitive basis for linguistic structures*. In J. Hayes (Ed.), *Cognition and the development of language* (pp. 279-362). Wiley.
- [3] Bisk, Y., Holtzman, A., & Thomason, J., et al. (2020). *Experience grounds language*. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP).
- [4] Caucheteux, C., & King, J. R. (2022). Brains and algorithms partially converge in natural language processing. *Communications Biology*, 5(1), 134.
- [5] Frazier, L., & Fodor, J. D. (1978). The sausage machine: A new two-stage parsing model. *Cognition*, 6(4), 291-325.
- [6] Hale, J. (2014). *Automaton theories of human sentence comprehension*. CSLI Publications.
- [7] Hu, J., & Levy, R. (2023). *A systematic assessment of syntactic generalization in neural language models*. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL 2023) (pp. 12333-12350).
- [8] Jain, S., & Huth, A. G. (2018). Incorporating context into language encoding models for fMRI. *Nature Neuroscience*, 21(4), 817-826.
- [9] Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review*, 99(1), 122-149.
- [10] Li, X., O'Sullivan, J. M., & Mattingley, J. B. (2022). Delay activity during visual working memory: A meta-analysis of 30 fMRI experiments. *NeuroImage*, 255, 119204. (<https://doi.org/10.1016/j.neuroimage.2022.119319>)
- [11] Linzen, T. (2020). *How can we accelerate progress towards human-like linguistic generalization?* In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL 2020) (pp. 5210-5217).
- [12] Marvin, R., & Linzen, T. (2018). *Targeted syntactic evaluation of language models*. In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP).
- [13] Merx, D., & Frank, S. L. (2021). *Comparing transformers and RNNs on predicting human sentence processing data*. In Proceedings of the CoNLL 2021 Shared Task: Cross-Framework Meaning Representation Parsing (pp. 74-86).
- [14] Misra, K. (2022). *When does a model become a subject? Scaling cognitive capabilities of language models*. In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP 2022) (pp. 1-15).
- [15] Pandia, L., & Ettinger, A. (2021). *Sorting through the noise: Testing sensitivity to hierarchical structure in neural language models*. In Proceedings of the 25th Conference on Computational Natural Language Learning (CoNLL 2021) (pp. 294-305).
- [16] Schrimpf, M., Blank, I. A., & Tuckute, G., et al. (2021). The neural architecture of language: Integrative modeling converges on predictive processing. *Proceedings of the National Academy of Sciences*, 118(45), e2105646118.
- [17] Taoudi-Benchekroun, Y., Christiaens, & D., Grigorescu, I., et al. (2022). Predicting age and clinical risk from the neonatal connectome. *NeuroImage*, 257, 119319. (<https://doi.org/10.1016/j.neuroimage.2022.119319>)
- [18] Tuckute, G., Sathe, A., & Srikant, S., et al. (2024). *Driving and suppressing the human language network using large language models*. *Nature Human Behaviour*.
- [19] Wehbe, L., Murphy, B., & Talukdar, P., et al. (2014). Simultaneously uncovering the patterns of brain regions involved in different story reading subprocesses. *PLOS Computational Biology*, 10(7), e1003930.