

Article

Not peer-reviewed version

Slang Context-Based Inference Enhancement via Greedy Search- Guided Chain-of-Thought Prompting

[Jinghan Cao](#)^{*}, Qingyang Ren, Xiangyun Chen, [Xinjin Li](#), Haoxiang Gao, Yu Zhao

Posted Date: 16 January 2026

doi: 10.20944/preprints202601.1159.v1

Keywords: slang inference; Large Language Model; chain-of-thought (CoT) prompting; semantic adaptation and understanding; natural language processing; foundational models for big data; machine learning



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Slang Context-Based Inference Enhancement via Greedy Search-Guided Chain-of-Thought Prompting

Jinghan Cao ^{1,*}, Qingyang Ren ², Xiangyun Chen ³, Xinjin Li ⁴, Haoxiang Gao ⁵ and Yu Zhao ⁶

¹ Department of Computer Science, San Francisco State University, San Francisco, USA

² Department of Computer Science, Cornell University, Ithaca, USA

³ Department of Biochemistry and Molecular Biology, Pennsylvania State University, University Park, USA

⁴ Department of Computer Science, Columbia University, New York, USA

⁵ Department of Electrical and Computer Engineering, Carnegie Mellon University, Mountain View, USA

⁶ Independent Researcher, New York, USA

* Correspondence: jcao3@alumni.sfsu.edu

Abstract

Slang interpretation has been a challenging downstream task for Large Language Models (LLMs) as the expressions are inherently embedded in contextual, cultural, and linguistic frameworks. In the absence of domain-specific training data, it is difficult for LLMs to accurately interpret slang meaning based on lexical information. This paper attempts to investigate the challenges of slang inference using large LLMs and presents a greedy search-guided chain-of-thought framework for slang interpretation. Through our experiments, we conclude that the model size and temperature settings have limited impact on inference accuracy. Transformer-based models with larger active parameters do not generate higher accuracy than smaller models. Based on the results of the above empirical study, we integrate greedy search algorithms with chain-of-thought prompting for small language models to build a framework that improves the accuracy of slang interpretation. The experimental results indicate that our proposed framework demonstrates improved accuracy in slang meaning interpretation. These findings contribute to the understanding of context dependency in language models and provide a practical solution for enhancing slang comprehension through a structured reasoning prompting framework.

Keywords: slang inference; Large Language Model; chain-of-thought (CoT) prompting; semantic adaptation and understanding; natural language processing; foundational models for big data; machine learning

I. Introduction

Large Language Models (LLMs) have been extensively investigated and applied across various downstream tasks, including content writing [1], summarization [2], information extraction [3], content understanding [4], and classification [5]. Trained on vast amounts of textual data, these models demonstrate proficiency in capturing complex language patterns and contextual relationships.

However, interpretation of informal languages, particularly slang expressions remains a significant challenge. While language models can easily infer the meaning if the slang datasets are used for pre-training, they may fail to correctly infer the meaning of unseen slang, even when context is provided. The rapid evolution of slang through social media and generational changes coupled with the substantial costs of model training, makes it impractical to continuously update models with emerging slang data. In this paper, we propose an enhanced chain-of-thought prompting approach that incorporates greedy search algorithms to improve models' reasoning capabilities for slang interpretation and meaning inference.

This study comprises two major components. First, we conducted a comparative analysis across different models to investigate the impact of temperature settings and model size on slang interpretation accuracy. Based on this empirical finding, we then proposed a greedy search-guided chain-of-thought framework and evaluated its performance in comparison with the conventional Input-Output (IO) prompting approach.

The primary contributions of this paper are as follows:

- We empirically evaluate two meta parameters (model size and temperature) in prompting approach to perform slang meaning inference. Our experimental results demonstrate that increasing model size yields limited improvements, while higher temperature settings may adversely affect task accuracy.
- We propose a novel application of chain-of-thought approach integrated with greedy search algorithm to enhance language models' inference capabilities. This approach extends the application of chain-of-thought prompting beyond its traditional use in arithmetic reasoning tasks, demonstrating its effectiveness in improving models' general inference abilities.

II. Related Work

Slang, as a form of informal language, has been rapidly evolving and adapting to the widespread social media and digital communication platforms [6]. However, in the field of natural language processing, slang-related research is still insufficient. Unlike mathematical reasoning tasks that converge to definitive solutions, slang interpretation is inherently dynamic and context-dependent, varying across cultural, linguistic, and regional dimensions [7–9]. Language models trained on historical corpus data have limited inference capability to “guess” a given slang term's primary meaning. For example, as shown in Figure 1, the term *Dope* has multiple meanings in different contexts. It can be used to express *something is pretty cool*, but it also refers to *stupid person* [10]. Given the dynamic nature of slang expressions, existing research can be classified into two major areas:

Slang Detection. To explore if the trained models are capable of identifying slang terms in everyday communication, researchers have made efforts to propose novel model training strategies using various data sources and to introduce benchmarks to assess models' capability in performing this task. For example, Sun et al. [11] constructed a slang dataset composed of dialogues extracted from English movies as the benchmark and enhanced the LLMs' ability to detect slang expressions in a sentence and infer their sources by applying pretrained Transformer models and fine-tuned GPT models. In addition to the innovation in the comparative analysis on models, Pei et al. [12] proposed a novel model structure by building bidirectional long short-term memory (Bi-LSTM) model with multilayer perceptron [13] to detect if slang exists at both sentence-level and token-level. Similarly, Seki and Liu [14] introduced a two-layered bidirectional LSTM with a hierarchical multi-task learning approach to detect Internet slang in Japanese.

Slang Interpretation. As early as 2017, Ni and Wang [15] collected large amounts of slang data from Urban Dictionary and utilized a dual encoder structure in an LSTM model to generate slang explanations. With the emergence of LLMs, Wuraola et al. [16] conducted research on employing transformer-based model in identifying the hidden emotional nuances from climate-related tweets and highlighted the shortcomings of language models in cultural insensitivity. Similarly, Mei et al. [17] extended the context of slang interpretation to daily communication scenarios and introduced a benchmark to evaluate existing language models' comprehension of new phrases given the colloquial context.

By reviewing the relevant research on slang detection and interpretation, a notable gap exists in approaches to augment a language model's intricate comprehension ability in performing downstream tasks. In this study, as illustrated in Figure 2, we propose a novel chain-of-thought prompting approach enhanced by greedy decoding algorithms to generate optimal context-aware explanations for novel slang expressions.

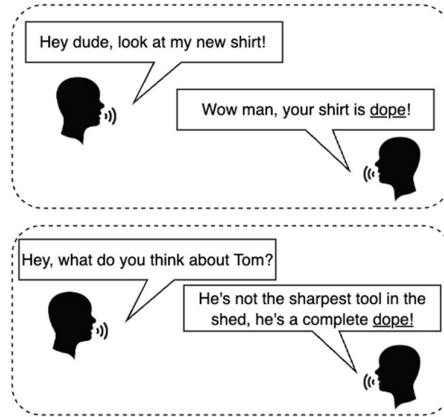


Figure 1. Context-Dependent Slang Interpretation.

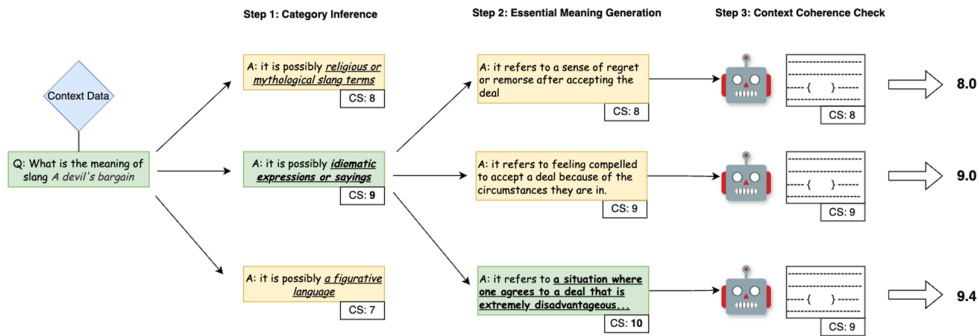


Figure 2. Greedy Search-Guided Chain-of-Thought.

III. Methodology

A. Problem Formulation

Each slang record consists of word S_{word} , ground truth meaning M_{origin} and usage example $C_{example}$. The objective is to generate the most-relevant slang meaning M_{gen} that has the highest semantic similarity with M_{origin} . Formally we break down this process into two main steps.

Firstly, given the input S_{word} and $C_{example}$, the best guessed slang meaning M_{gen} is generated as the output

$$M_{gen} = f(S_{word}, C_{example}; P_{cot}) \quad (1)$$

where P_{cot} represents the chain-of-thought prompts.

Secondly, given input x as an element of M_{gen} and the ground truth meaning M_{origin} , the output S represents the generated meaning with the highest semantic similarity score relative to M_{origin}

$$S = \operatorname{argmax}_{x \in M_{gen}} \operatorname{sim}(x, M_{origin}) \quad (2)$$

where $\operatorname{sim}(x, M_{origin})$ denotes the semantic similarity between generated meaning and the ground truth meaning.

B. Greedy Search-Guided Chain-of-Thought

Chain-of-thought strategy has been proven effective in improving the ability of LLMs to perform complex reasoning tasks [18]. Inspired by the self-consistent chain-of-thought strategy [19] and tree-of-thought strategy [20], we have developed a specialized chain-of-thought workflow for inferring slang meanings. As indicated in Figure 3, we introduce confidence score and employ a greedy algorithm to identify the most probable explanation for slang expressions. Unlike the conventional

Input-Output (IO) single chain, our method expands the number of candidate thoughts to W with a depth of D , where W represents the number of possible candidates at each step and D represents the total number of steps in the chain. Both W and D are task-dependent parameters. In our experiments, we set both W and D to 3, aligning with the designed prompting steps.




Figure 3. Greedy Algorithm in Chain-of-Thought.

C. Step-by-Step Prompting

We break the chain-of-thinking process into a three-stage inference framework.

Category. An initial prompt is designed to help the model identify possible categories of slang based on given usage examples. As described in Algorithm 1 following principal of greedy algorithm, we select the inferred category with highest confidence score.

Algorithm 1 Category Inference

Inputs: *Context*: the usage example; *Slang*: the target slang term; K : total number of candidates to be generated.

Prompt: prompt to instruct LLM to generate categories.

Output: *maxTuple*: the category with the highest score

```

1: maxTuple  $\leftarrow$  (null, 0)
2: thoughts  $\leftarrow$  LLMInfer (Context, Slang,  $K$ , Prompt)
3: for (category, score) in thoughts do
4:   if maxTuple [1] < score then
5:     maxTuple  $\leftarrow$  (category, score)
6:   end if
7: end for
8: return maxTuple

```

Essential Meaning. In the second prompting iteration, we instruct the LLM to generate possible primary meanings based on the previously inferred category. The process is similar to Algorithm 1 except for an additional input parameter representing the inferred category. The output is a list of candidate meanings for the target slang.

Algorithm 2 Essential Meaning Generation

Inputs: *Category*: the selected category; *Slang*: the target slang term; K : total number of candidates to be generated. *Prompt*: prompt to instruct LLM to generate meanings.

Output: *meaningList*: a list of generated meanings with confidence score

```

1: meaningList  $\leftarrow$  []
2: meaningList  $\leftarrow$  LLMInfer (Category, Slang,  $K$ , Prompt)
3: return meaningList

```

Compatibility. In the final step, we select the slang interpretation with the highest confidence score and prompt the LLM to evaluate whether this meaning aligns with the original context and tone. A heuristic weighting mechanism is applied in calculating the final confidence score.

Algorithm 3 Context Coherence Check

Inputs: *Context*: the usage example; *meaningList*: a list of generated meanings with confidence score. *Prompt*: prompt to instruct LLM to check compatibility in original context

Output: *selectedMeaning*: the selected meaning

```

1: selectedMeaning ← null
2: finalScore ← 0
3: selectedMeaning ← null
4: for (meaning, score) in meaningList do
5:   confidenceScore ← LLMInfer (Context, meaning, Prompt)
6:   if finalScore < (confidenceScore * 0.6 + score * 0.4)
7:     finalScore ← confidenceScore * 0.6 + score * 0.4
8:     selectedMeaning ← meaning
9:   end if
10: end for
11: return selectedMeaning

```

IV. Experiment and Analysis

We conducted two sets of experiments. In the first experiment, we evaluated the inference capabilities across different model sizes and temperature settings using standard IO prompting. Based on these findings, we then proceed with the second experiment implementing the proposed greedy chain-of-thought prompting strategy and comparing its performance with the conventional CoT approach.

A. Experimental Setting

Configuration. In the first experiment, we evaluated multiple language models, including GTP-4o, Qwen2.5-72B, DeepSeek-V3 and their smaller versions including GPT-4o-mini, Qwen2-7B-Instruct and DeepSeek-R1-Distill-Llama-8B. For evaluation, we randomly selected and preprocessed 1,200 high-quality slang records for testing across GPT, Qwen, and DeepSeek models. Additionally, we investigated the impact of temperature settings using Qwen2-7B-Instruct as the baseline model, testing temperatures of 0.1, 0.3, 0.5, and 0.7 across 500 slang records.

In the second experiment, we compared standard IO promoting with our proposed chain-of-thought approach incorporating greedy search using Qwen2-7B-Instruct as the base model. The temperature was set to 0.3 and tested against 1200 random processed slang records.

For both experiments, we utilized OpenAI APIs for GPT series models, while Qwen and DeepSeek models were locally hosted and evaluated on a cluster of 8 NVIDIA A100 GPUs with a total of 320GB GPU memory.

Dataset Preprocessing. The raw dataset comprises highly-voted slang entries extracted from Urban Dictionary [10], with each record containing three attributes: slang word, meaning, and usage example. Given the variable quality of Urban Dictionary records, we employed GPT-4o to filter out mismatched entries and standardize both ground truth meanings and usage examples. For usage examples specifically, we reformatted the raw content into structured dialogues to provide clearer context (Figure 4).

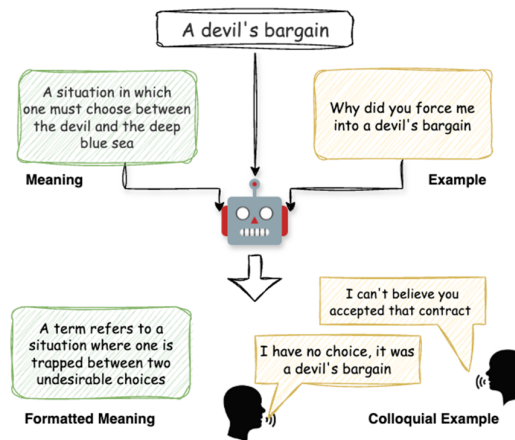


Figure 4. Rephrase the original meaning and example.

B. Metrics

ROUGE-L. To evaluate the similarity between inferred meanings and ground-truth references, we employ ROUGE-L [21], which computes sentence-level content similarity through F1 score, precision, and recall metrics.

SimCSE. This sentence-level embedding benchmark [22] demonstrates superior performance in semantic textual similarity tasks using supervised BERT-based models.

C. Experimental Results

- a) *Limited Impact from Model Size and Temperature Setting:* Based on experimental results comparing multiple model sizes, we observe that larger models do not necessarily demonstrate superior performance, as shown in Table 1. Interestingly, smaller models such as GPT-4o-mini and Qwen2-7B-Instruct achieve higher F1 scores despite lower SimCSE results.

Table 1. Impact of Model Size on Inference.

Models	ROUGE-L			SimCSE
	F1	Precision	Recall	
GPT-4o	0.225	0.149	0.171	0.736
GTP-4o-mini	0.299	0.166	0.250	0.727
Qwen2.5-72B	0.123	0.250	0.199	0.715
Qwen2-7B-Instruct	0.170	0.322	0.222	0.696
DeepSeek-V3	0.235	0.166	0.300	0.726
DeepSeek-R1-Distill-Llama-8B	0.166	0.111	0.222	0.696

Then we conducted experiments with incremental temperature settings on Qwen2-7B-Instruct and DeepSeek-R1-Distill-Llama-8B. As shown in Figure 5, our experiments reveal that higher temperature settings do not correlate with improved slang inference performance. This finding aligns with prior research [23], which established that temperature variations have minimal impact on LLMs' problem-solving capabilities.





Figure 5. Temperature Impact on Inference Ability.



- b) *Improved Accuracy with Greedy Search-Guided Chain-of-Thought:* Given the superior overall performance of Qwen2-7B-Instruct with the temperature set to 0.3, we proceeded to the second experiment implementing our proposed approach.

As illustrated in Figure 6, the comparative analysis revealed significant performance differences. The standard Input-Output approach achieved an F1 score of 0.235, precision of 0.261, recall of 0.242, and SimCSE score of 0.696. In contrast, our proposed greedy chain-of-thought method demonstrated superior performance across all metrics, achieving an F1 score of 0.500, precision of 0.363, recall of 0.411, and SimCSE score of 0.704.




Figure 6. Performance Comparison between standard IO Prompt and Greedy Search-Guided CoT Prompt.

The experimental results highlight the effectiveness of greedy search-guided chain-of-thought prompts. Unlike traditional reasoning tasks, slang inference presents unique challenges in evaluating generated candidates at each step, as the LLM lacks access to the ground-truth meaning. Our approach addresses this by expanding the number of generated candidates heuristically, setting the candidate count to three for initial performance verification.

Analysis of experiment logs revealed that the LLM's first response is not consistently the most confident. This can be attributed to the non-deterministic nature of LLMs and the complex token sampling mechanisms that generate varying results. Moreover, confidence score assignment operates as an independent process, influenced by input parameters and meta-configuration settings.

V. Conclusions

This paper presents a novel approach to enhance slang interpretation capabilities in language models through a greedy search-guided chain-of-thought prompting framework. Our research reveals several significant findings:

First, our empirical analysis revealed that model size and temperature settings have limited impact on slang interpretation accuracy. Contrary to common assumptions, larger models did not consistently outperform their smaller counterparts, and higher temperature settings did not lead to improved inference capabilities.

Second, our proposed framework demonstrated substantial improvements over standard Input-Output prompting, achieving higher scores across all evaluation metrics. This improvement can be attributed to the framework's structured approach to reasoning and its ability to evaluate multiple candidate interpretations at each step.

The success of our approach suggests that breaking down slang interpretation into discrete reasoning steps enables more accurate and contextually appropriate interpretations. This finding has important implications for improving language models' handling of informal language without requiring additional training or model parameter modifications.

Future work will explore generalization of this approach to be applied in other inference tasks and develop more sophisticated confidence scoring mechanisms for candidate evaluation.

Appendix A

This section outlines the step-by-step sample prompts to instruct the model to generate expected results.

```

System Prompt
You are an expert in urban slang and internet language

User Prompt
Given the context:
Example usage: (context)
Please explain the following word:
Word: (slang_word)
Based on the word morphology and Example usage, infer which category this slang term most likely belongs to, provide 3 possible options with confidence score. You must return a JSON response with exactly this format
{
  "Your_First_Thought": "your first thought"
  "Your_First_Thought_Score": [0-10]
  "Your_Second_Thought": "your second thought"
  "Your_Second_Thought_Score": [0-10]
  "Your_Third_Thought": "your third thought"
  "Your_Third_Thought_Score": [0-10]
}

```

Figure A1. Sample Prompt to Infer Category.

The first step in the prompt chain requires the model to infer the top-K ($K=3$) possible categories before generating any specific meanings.

```

System Prompt
You are an expert in urban slang and internet language

User Prompt
Given the original context:
Example usage: (original_context)
And the most likely inferred category:
Inferred category: (inferred_category)
Based on Inferred category, what is the essential meaning of this term? Provide 3 possible options with confidence score.
You must return a JSON response with exactly this format
{
  "Your_First_Thought": "your first thought"
  "Your_First_Thought_Score": [0-10]
  "Your_Second_Thought": "your second thought"
  "Your_Second_Thought_Score": [0-10]
  "Your_Third_Thought": "your third thought"
  "Your_Third_Thought_Score": [0-10]
}

```

Figure A2. Sample Prompt to Generate Possible Meanings.

The LLM takes the most likely category as input for the next prompt. Based on this input, it generates three possible meanings and assigns confidence scores to each.

```

System Prompt
You are an expert in urban slang and internet language

User Prompt
Given the original context:
Example usage: (original_context)
And the most likely inferred category:
Inferred category: (inferred_category)
And the most likely inferred meaning:
Inferred meaning: (inferred_meaning)
Evaluate if the inferred meaning fit into the original context and provide your confidence score in terms of the compatibility.
{
  "Your_Confidence_Score": [0-10]
}

```

Figure A3. Sample Prompt to Provide Confidence Score for Selected Meaning.

In the final step, the model evaluates the compatibility between the original context, inferred categories, and derived meanings to generate a confidence score.

References

1. Y. Zeng et al., 'Bridging the Editing Gap in LLMs: FineEdit for Precise and Targeted Text Modifications', arXiv [cs.CL]. 2025.
2. Y. Ji et al., 'RAG-RLRC-LaySum at BioLaySumm: Integrating Retrieval-Augmented Generation and Readability Control for Layman Summarization of Biomedical Texts', in Proceedings of the 23rd Workshop on Biomedical Natural Language Processing, 2024, pp. 810–817.
3. Y. Ji, Z. Yu, and Y. Wang, 'Assertion Detection in Clinical Natural Language Processing Using Large Language Models', in 2024 IEEE 12th International Conference on Healthcare Informatics (ICHI), 2024, pp. 242–247.
4. X. Lin, Y. Tu, Q. Lu, J. Cao, H. Yang, and Others, 'Research on Content Detection Algorithms and Bypass Mechanisms for Large Language Models', Academic Journal of Computing & Information Science, vol. 8, no. 1, pp. 48–56.
5. C. Shi et al., 'Deep Semantic Graph Learning via LLM based Node Enhancement', arXiv [cs.AI]. 2025.
6. A. Sundaram, H. Subramaniam, S. H. A. Hamid, and A. M. Nor, "A Systematic Literature Review on Social Media Slang Analytics in Contemporary Discourse", IEEE Access, vol. 11, pp. 132457–132471, 2023.
7. "Slang and Sociability | Connie Eble," University of North Carolina Press. <https://uncpress.org/book/9780807845844/slang-and-sociability/>
8. D. Hovy and A. Søgaard, "Tagging Performance Correlates with Author Age", in Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers), 2015, pp. 483–488.
9. D. Hovy and S. L. Spruit, "The Social Impact of Natural Language Processing", in Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), 2016, pp. 591–598.
10. Urban Dictionary, "Urban Dictionary," Urban Dictionary, 2000. <https://www.urbandictionary.com/>
11. Z. Sun, Q. Hu, R. Gupta, R. Zemel, and Y. Xu, "Toward Informal Language Processing: Knowledge of Slang in Large Language Models", in Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers), 2024, pp. 1683–1701.
12. Z. Pei, Z. Sun, and Y. Xu, "Slang Detection and Identification", in Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL), 2019, pp. 881–889.
13. T. Rauber and K. Berns, "Kernel Multilayer Perceptron", in 2011 24th SIBGRAPI Conference on Graphics, Patterns and Images, 2011, pp. 337–343.
14. Y. Seki and Y. Liu, 'Multi-task Learning Model for Detecting Internet Slang Words with Two-Layer Annotation', in 2022 International Conference on Asian Language Processing (IALP), 2022, pp. 212–218.
15. K. Ni and W. Y. Wang, "Learning to Explain Non-Standard English Words and Phrases", in Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers), 2017, pp. 413–417.
16. I. Wuraola, N. Dethlefs, and D. Marciniak, 'Understanding Slang with LLMs: Modelling Cross-Cultural Nuances through Paraphrasing', in Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, 2024, pp. 15525–15531.
17. L. Mei, S. Liu, Y. Wang, B. Bi, and X. Cheng, "SLANG: New Concept Comprehension of Large Language Models", in Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, 2024, pp. 12558–12575.
18. J. Wei et al., 'Chain-of-Thought Prompting Elicits Reasoning in Large Language Models', arXiv [cs.CL]. 2023.
19. X. Wang et al., 'Self-Consistency Improves Chain of Thought Reasoning in Language Models', arXiv [cs.CL]. 2023.

20. S. Yao et al., 'Tree of Thoughts: Deliberate Problem Solving with Large Language Models', arXiv [cs.CL]. 2023.
21. C.-Y. Lin, 'ROUGE: A Package for Automatic Evaluation of Summaries', in Text Summarization Branches Out, 2004, pp. 74–81.
22. T. Gao, X. Yao, and D. Chen, 'SimCSE: Simple Contrastive Learning of Sentence Embeddings', arXiv [cs.CL]. 2022.
23. M. Renze, 'The Effect of Sampling Temperature on Problem Solving in Large Language Models', in Findings of the Association for Computational Linguistics: EMNLP 2024, 2024, pp. 7346–7356.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.