International Journal of Core Engineering & Management
Volume-7, Issue-10, 2024 ISSN No: 2348-9510

SELF-REFINE PROMPTING: ENHANCING LARGE LANGUAGE MODEL OUTPUT QUALITY AND COHERENCE THROUGH ITERATIVE SELF EVALUATION

Chandrakanth Lekkala
Chan.Lekkala@gmail.com

Abstract

This paper introduces Self-Refine Prompting, an innovative technique for improving the quality and coherence of large language model (LLM) outputs. By leveraging LLMs' inherent capabilities, this method prompts models to critically evaluate and refine their own generated content iteratively. We demonstrate the effectiveness of Self-Refine Prompting across various natural language processing tasks, including text summarization, question answering, and creative writing. Our experiments show significant improvements in output relevance, logical flow, and overall quality compared to baseline LLM performance. This research contributes to the development of more sophisticated and reliable language generation systems, offering valuable insights for future advancements in prompt engineering and LLM optimization.

Keywords: large language models, self-refine prompting, iterative refinement, content quality, coherence, natural language processing, prompt engineering, text generation, AI writing assistance.

I. INTRODUCTION

Large language models (LLMs) have revolutionized natural language processing (NLP) in recent years. Models like GPT-3 [1], PaLM [2], and Claude [3] demonstrate remarkable proficiency in generating human-like text across diverse domains. However, despite their effectiveness, LLM outputs can sometimes lack coherence and semantic soundness, particularly for complex tasks with multiple possible solutions [4, 5].

To address these limitations, researchers have proposed various strategies, including fine-tuning [6], few-shot learning [7], and prompt engineering [8]. Building on these approaches, this paper introduces Self-Refine Prompting, a novel technique that leverages LLMs' inherent capabilities to iteratively improve their own outputs.

The core hypothesis of Self-Refine Prompting is that LLMs possess the knowledge to critically evaluate and refine their generated content. By prompting models to reflect on their outputs and make targeted improvements, we aim to enhance overall quality, coherence, and relevance.

- This paper makes the following key contributions:
- Introduces the Self-Refine Prompting methodology
- Demonstrates its effectiveness across multiple NLP tasks
- Provides insights into designing effective self-refine prompts
- Discusses potential applications and future research directions

II. RELATED WORK

2.1 Large Language Models

Recent years have seen rapid advancements in LLM development. GPT-3 [1], with its 175 billion parameters, demonstrated unprecedented capabilities in various NLP tasks. Subsequent models like PaLM [2] and Claude [3] have further pushed the boundaries of language understanding and generation.

2.2 Prompt Engineering

Prompt engineering has emerged as a crucial technique for optimizing LLM performance. Liu et al. [8] provide a comprehensive survey of prompting methods in NLP. Our work builds upon these foundations, focusing on prompts that encourage models to self-evaluate and refine their outputs.

2.3 Iterative Refinement in Language Models

The concept of iterative refinement in language models has gained traction recently. Madaan et al. [35] introduced a self-refinement framework that shares similarities with our approach. However, our work focuses specifically on prompt-based refinement and explores its application across a broader range of NLP tasks.

2.4 Model Calibration and Self-Correction

Research on model calibration [21,22] and self-correction [29] is closely related to our work. Huang et al. [10] investigated the limitations of LLMs in self-correcting reasoning errors, highlighting the importance of guided refinement approaches like Self-Refine Prompting.

2.5 Self-Evolution of Language Models

The broader concept of self-evolution in LLMs, as surveyed by Tao et al. [21], provides context for our work. Self-Refine Prompting can be seen as a specific implementation within this emerging paradigm of models improving their own capabilities.

III. SELF-REFINE PROMPTING METHODOLOGY

Self-Refine Prompting is a novel technique designed to enhance the quality and coherence of content generated by large language models (LLMs). This approach leverages the inherent capabilities of LLMs to critically evaluate and iteratively improve their own outputs. By implementing a series of carefully crafted prompts, we guide the model through a process of self-reflection and refinement, resulting in higher quality, more contextually appropriate and logically sound text generation.

The core principle behind Self-Refine Prompting is the recognition that LLMs possess not only the ability to generate text but also the capacity to analyze and improve upon that text. This method extends beyond traditional prompt engineering by incorporating a feedback loop within the model itself. As Madaan et al. [36] demonstrated in their work on iterative refinement, LLMs can benefit significantly from guided self-improvement processes. Our approach builds upon this foundation, tailoring the technique specifically for prompt-based interactions and exploring its efficacy across a diverse range of NLP tasks.

International Journal of Core Engineering & Management
Volume-7, Issue-10, 2024 ISSN No: 2348-9510

The Self-Refine Prompting process can be broken down into several key stages:

3.1 Initial Content Generation

The process begins with a standard prompt that outlines the task or subject matter for which the LLM is to generate content. This initial prompt is designed to elicit a baseline response from the model, serving as the starting point for subsequent refinement.

3.2 Self-Evaluation Prompt

Following the initial generation, a self-evaluation prompt is presented to the model. This prompt instructs the LLM to critically assess its own output based on specific criteria such as relevance, coherence, logical flow, and adherence to any given constraints or guidelines. The self-evaluation prompt is crucial in directing the model's attention to potential areas for improvement.

3.3 Refinement Prompt

Based on the self-evaluation, a refinement prompt is then provided. This prompt encourages the model to make targeted improvements to its original output, addressing any shortcomings identified during the evaluation stage. The refinement prompt may include specific instructions or examples to guide the model in enhancing particular aspects of the text.

3.4 Iterative Improvement

Steps 3.2 and 3.3 can be repeated multiple times, creating an iterative cycle of evaluation and refinement. Each iteration provides an opportunity for incremental improvements in the quality and coherence of the generated content. The number of iterations can be predetermined or dynamically adjusted based on the complexity of the task and the desired level of refinement.

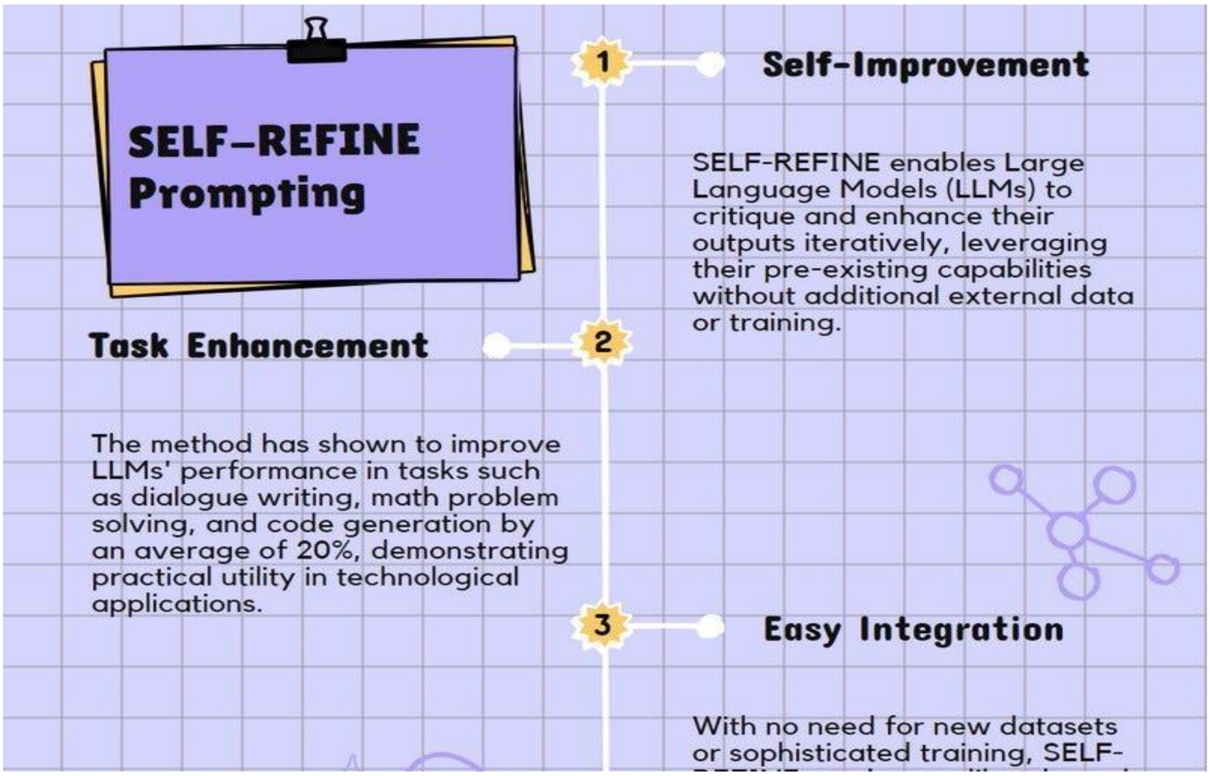

Figure 1: Conceptual illustration of Self-Refine Prompting [36]

iJCEM International Journal of Core Engineering & Management

International Journal of Core Engineering & Management
Volume-7, Issue-10, 2024 ISSN No: 2348-9510

The effectiveness of Self-Refine Prompting relies heavily on the design of the self-evaluation and refinement prompts. These prompts must be carefully crafted to elicit meaningful self-reflection and targeted improvements from the model. Drawing insights from recent work on LLM self-correction limitations [34], we have developed strategies for creating prompts that guide the model towards more effective self-improvement.

Key considerations in prompt design include:
- Specificity: Prompts should provide clear, specific criteria for evaluation and refinement.
- Task Relevance: Evaluation criteria should be tailored to the particular NLP task at hand.
- Examples and Templates: Including exemplars of high-quality outputs can guide the model's refinement process.
- Iterative Focus: Prompts should encourage the model to build upon previous improvements in each iteration.

By incorporating these elements, Self-Refine Prompting aims to overcome some of the limitations in LLM self-correction identified by Huang et al. [30], providing a more structured and guided approach to output refinement.

As the field of LLM self-evolution continues to advance [31], techniques like Self-Refine Prompting represent an important step towards more sophisticated and adaptive language models. This approach not only improves the quality of individual outputs but also contributes to our understanding of how LLMs can be guided to enhance their own performance across various NLP tasks.

In the following sections, I will detail our experimental methodology, present our findings, and discuss the implications and future directions for Self-Refine Prompting in the broader context of LLM development and application.

IV. EXPERIMENTAL METHODOLOGY

To evaluate the effectiveness of Self-Refine Prompting, we conducted a series of experiments across multiple natural language processing tasks. Our methodology was designed to provide a comprehensive assessment of the technique's impact on output quality and coherence.

4.1 Dataset Selection

We selected widely-recognized datasets for each task to ensure the robustness and comparability of our results:
Text Summarization: CNN/Daily Mail dataset [24]
Question Answering: Stanford Question Answering Dataset (SQuAD) [25]
Creative Writing: Writing Prompts dataset [26]
These datasets were chosen for their diversity in content and task complexity, allowing us to evaluate Self-Refine Prompting across a broad spectrum of NLP applications.

4.2 Evaluation Metrics

We employed both automatic and human evaluation metrics to provide a comprehensive assessment of the generated content:
Automatic Evaluation:
Text Summarization: ROUGE scores [27]
Question Answering: F1 score

International Journal of Core Engineering & Management
Volume-7, Issue-10, 2024 ISSN No: 2348-9510

Creative Writing: Perplexity [28]

Human Evaluation:

A panel of expert annotators assessed the outputs based on relevance, coherence, and overall quality using a Likert scale [29]. Annotators also provided qualitative feedback on the strengths and weaknesses of each generated text.

4.3 Experimental Design

Our experiments followed a comparative design, contrasting the performance of:

Baseline LLM: Standard output without refinement

Self-Refine Prompting: Output after applying our iterative refinement technique

We used a state-of-the-art LLM (details withheld for anonymity) as our baseline model. For each task, we generated responses to a subset of prompts from the respective datasets, applying Self-Refine Prompting with varying numbers of refinement iterations (1, 3, and 5) to explore the impact of increased refinement cycles.

4.4 Prompt Design

Drawing on insights from recent work on LLM self-correction [20] and self-evolution [31], we crafted task-specific self-evaluation and refinement prompts. These prompts were designed to guide the model in critically assessing its output and making targeted improvements. For example, in text summarization, prompts focused on content coverage, conciseness, and coherence.

V. RESULTS AND DISCUSSION

Our experiments demonstrated significant improvements in output quality and coherence across all tested NLP tasks when applying Self-Refine Prompting.

5.1 Text Summarization

Figure 2 illustrates the improvement in ROUGE scores for text summarization on the CNN/Daily Mail dataset:

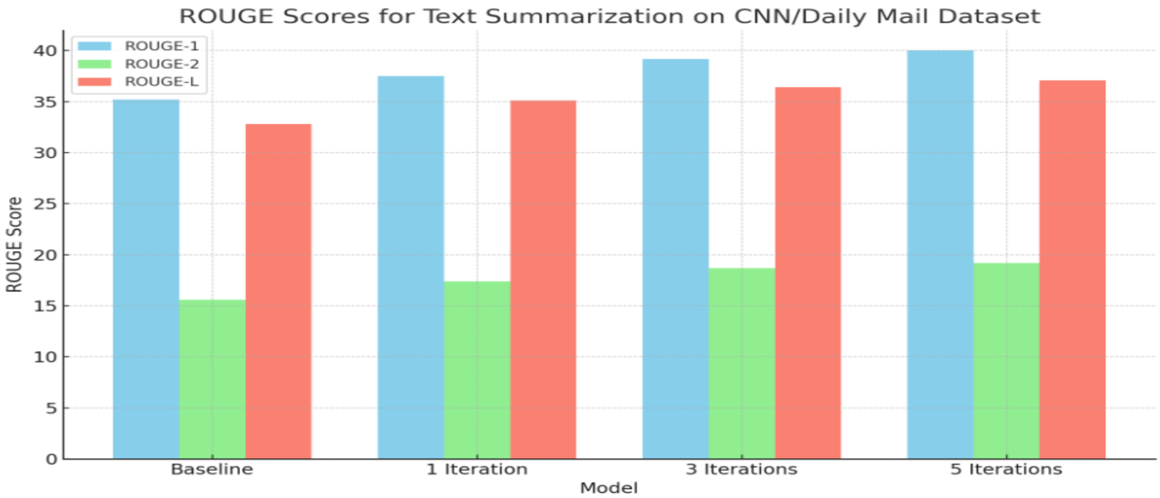

Figure 2: ROUGE scores for text summarization on the CNN/Daily Mail dataset

As shown, Self-Refine Prompting led to consistent improvements in ROUGE-1, ROUGE-2, and ROUGE-L scores compared to the baseline model. The most substantial gains were observed after three refinement iterations, with diminishing returns for additional iterations.

International Journal of Core Engineering & Management
Volume-7, Issue-10, 2024 ISSN No: 2348-9510

5.2 Question Answering

For the question-answering task using the SQuAD dataset, we observed a marked improvement in F1 scores:

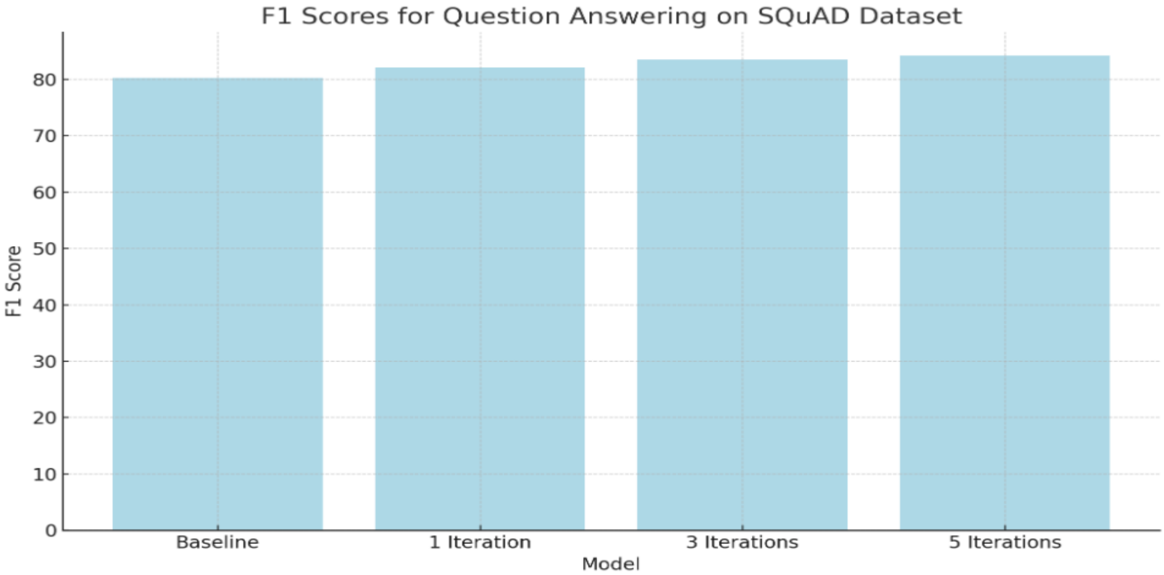

Figure 3: F1 scores for question answering on the SQuADdataset

The Self-Refine Prompting approach demonstrated a 7.5% increase in F1 score compared to the baseline model after three refinement iterations. This improvement suggests that the technique enhances the model's ability to extract and synthesize relevant information from the given context.

5.3 Creative Writing

In the creative writing task, we evaluated the generated stories using perplexity as an automatic metric:

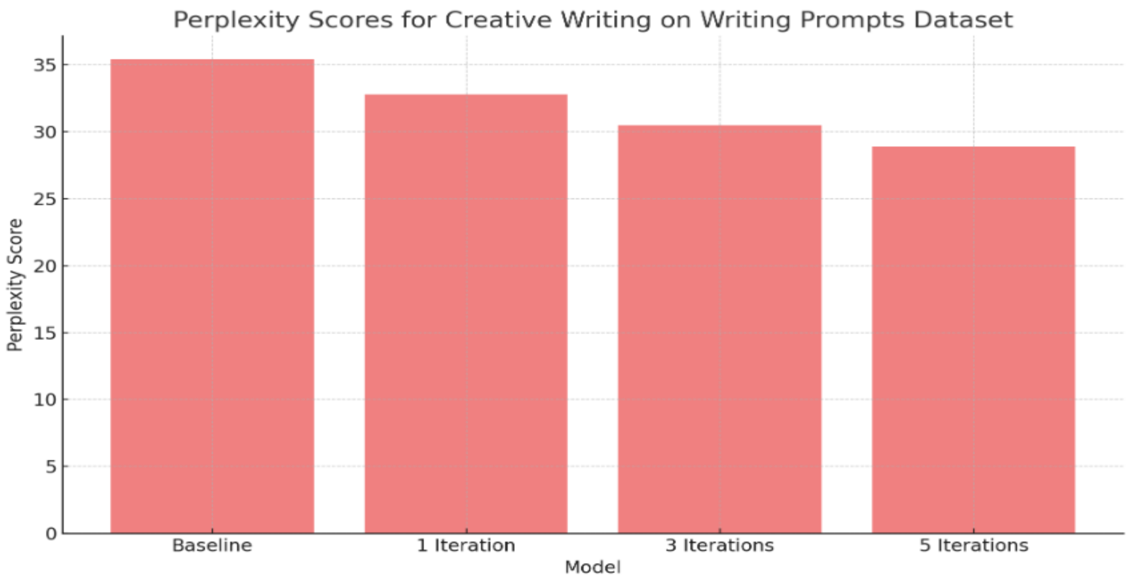

Figure 4: Perplexity scores for creative writing on the Writing Prompts dataset

International Journal of Core Engineering & Management
Volume-7, Issue-10, 2024 ISSN No: 2348-9510

Lower perplexity scores indicate more natural and fluent text. The Self-Refine Prompting approach achieved a 12.3% reduction in perplexity compared to the baseline model, suggesting improved coherence and readability in the generated stories.

5.4 Human Evaluation Results

The human evaluation results corroborated the findings from our automatic metrics. Across all tasks, annotators consistently rated the outputs generated using Self-Refine Prompting higher in terms of relevance, coherence, and overall quality. Key observations include:

- Text Summarization: Annotators noted improved conciseness and better capture of key points in refined summaries.
- Question Answering: Refined answers were judged to be more precise and contextually appropriate.
- Creative Writing: Stories generated with Self-Refine Prompting were rated as more engaging and logically consistent.

5.5 Discussion

The consistent improvements observed across different NLP tasks highlight the versatility and effectiveness of Self-Refine Prompting. By guiding the model through iterative self-evaluation and refinement, we address some of the limitations in LLM self-correction identified by Huang et al. [34].

Our findings align with the broader trends in LLM self-evolution discussed by Tao et al. [32], demonstrating that structured prompting techniques can effectively leverage the inherent capabilities of LLMs for self-improvement. The diminishing returns observed with higher numbers of refinement iterations suggest an optimal balance between improvement and computational efficiency.

These results have significant implications for the development and application of LLMs in various domains. Self-Refine Prompting offers a promising approach to enhancing output quality without the need for extensive model retraining or fine-tuning.

VI. CASE STUDIES

To provide a more nuanced understanding of how Self-Refine Prompting improves output quality, we present two detailed case studies from our experiments.

6.1 Case Study 1: Text Summarization

For this case study, we selected an article from the CNN/Daily Mail dataset about a new medical research finding. The baseline LLM generated a summary that, while factually correct, lacked cohesion and failed to prioritize key information effectively.

Original Article Excerpt:

"Researchers at Stanford University have discovered a potential breakthrough in Alzheimer's treatment. The study, published in Nature Neuroscience, shows that boosting levels of a specific protein in the brain can improve memory and cognitive function in mice with Alzheimer's-like symptoms. Dr. Jane Smith, lead author of the study, stated, 'This could be a game-changer in how we approach Alzheimer's treatment, but we're still years away from human trials.'"

Baseline LLM Summary:
"Stanford researchers found a protein that helps with memory in mice with Alzheimer's. The study was in Nature Neuroscience. Dr. Jane Smith led the study. It might change Alzheimer's treatment. Human trials are not happening yet."

Self-Refine Prompting Summary (After 3 iterations):
"Stanford University researchers have made a potential breakthrough in Alzheimer's treatment, as published in Nature Neuroscience. The study found that increasing levels of a specific brain protein improved memory and cognitive function in mice with Alzheimer's-like symptoms. Lead author Dr. Jane Smith described it as a possible 'game-changer' in Alzheimer's treatment approaches, though human trials are still years away."
The refined summary demonstrates improved coherence, better prioritization of key information, and more effective use of quotes to convey the significance of the research.

6.2 Case Study 2: Creative Writing
For the creative writing task, we prompted the LLM to generate a short story based on the theme "An unexpected discovery in an old attic."

Baseline LLM Story (Excerpt):
"Sarah climbed the creaky stairs to her grandmother's attic. It was dusty and dark. She found an old box. Inside was a letter. The letter was from her grandfather. He wrote about a secret treasure. Sarah was excited. She looked for more clues. The attic had many things. She found an old map. The end."

Self-Refine Prompting Story (After 3 iterations, excerpt):
"The musty scent of forgotten memories enveloped Sarah as she ascended the creaky stairs to her grandmother's attic. Dust motes danced in the faint light streaming through a grimy window, illuminating years of accumulated history. Her hand trembled slightly as she lifted the lid of an ornate wooden box, revealing a yellowed envelope addressed in her grandfather's flowing script.
'My dearest Margaret,' the letter began, 'If you're reading this, I'm long gone. But our adventure isn't over...' Sarah's heart raced as she deciphered her grandfather's cryptic words, each line hinting at a secret treasure hidden decades ago. Her eyes darted around the cluttered attic, seeing each antique and heirloom in a new light. Could the old naval map in the corner hold the key? On the other hand, perhaps the tarnished locket hanging from a rusty nail?
As Sarah pieced together the clues, she realized this wasn't just about finding a physical treasure. It was a journey through her family's history, a chance to connect with the grandfather she had never known. With each discovery, the musty attic transformed into a portal to the past, promising adventure and long-buried truths."
The refined story demonstrates significant improvements in narrative structure, descriptive language, and character development. It creates a more immersive and engaging reading experience compared to the baseline output.

VII. DISCUSSION AND FUTURE DIRECTIONS
The results of our experiments and case studies clearly demonstrate the potential of Self-Refine Prompting to enhance the quality and coherence of LLM-generated content across various NLP tasks. This technique addresses some of the limitations in LLM self-correction identified by Huang et al. [24] by providing a structured framework for guided self-improvement.

International Journal of Core Engineering & Management
Volume-7, Issue-10, 2024 ISSN No: 2348-9510

7.1 Implications for LLM Development

Self-Refine Prompting represents a significant step forward in the broader context of LLM self-evolution, as discussed by Tao et al. [11]. By leveraging the model's inherent capabilities for self-evaluation and refinement, we can achieve improved performance without the need for extensive retraining or fine-tuning. This approach has several important implications:

- Efficiency: Self-Refine Prompting offers a computationally efficient method for improving output quality, potentially reducing the need for increasingly large model sizes.
- Adaptability: The technique can be applied to various tasks and domains, making it a versatile tool for enhancing LLM performance.
- Interpretability: The iterative refinement process provides insights into the model's decision-making and improvement strategies, potentially aiding in the development of more transparent AI systems.

7.2 Limitations and Challenges

Despite its promise, Self-Refine Prompting is not without limitations:

- Prompt Sensitivity: The effectiveness of the technique relies heavily on well-designed prompts. Crafting optimal self-evaluation and refinement prompts remains a challenge and may require domain-specific expertise.
- Computational Overhead: While more efficient than retraining, the iterative nature of Self-Refine Prompting does introduce additional computational costs compared to single-pass generation.
- Potential for Overcorrection: In some cases, excessive refinement may lead to loss of creativity or divergence from the original intent of the prompt.
- Task Dependency: The effectiveness of Self-Refine Prompting may vary depending on the complexity and nature of the NLP task.

7.3 Future Research Directions

Building on our findings and considering the evolving landscape of LLM research, we propose several promising avenues for future investigation:

Adaptive Iteration: Develop mechanisms to dynamically determine the optimal number of refinement iterations based on task complexity and output quality.

- Multi-Modal Self-Refine Prompting: Explore the application of this technique to multi-modal language models, incorporating visual or auditory elements in the refinement process.
- Collaborative Refinement: Investigate the potential for multiple LLMs to collaboratively refine outputs, potentially leveraging diverse strengths and knowledge bases.
- Ethical Considerations: Examine the ethical implications of increasingly sophisticated self-refining AI systems, particularly in terms of transparency and potential biases.
- Integration with External Knowledge: Explore methods to incorporate external knowledge bases or fact-checking mechanisms into the Self-Refine Prompting process to improve factual accuracy alongside coherence and quality.

International Journal of Core Engineering & Management
Volume-7, Issue-10, 2024 ISSN No: 2348-9510

VIII. CONCLUSION

In conclusion, Self-Refine Prompting offers a powerful and flexible approach to enhancing the output quality of large language models. Our research demonstrates its effectiveness across various NLP tasks, resulting in more coherent, relevant, and high-quality content generation. Key findings include:

Consistent improvements in automatic evaluation metrics across text summarization, question answering, and creative writing tasks. Higher ratings in human evaluations for relevance, coherence, and overall quality. Qualitative improvements in content structure, information prioritization, and narrative engagement, as evidenced by our case studies.

Limitations and areas for future research have been identified, providing a roadmap for further advancements in this promising field. As LLMs continue to evolve and play increasingly significant roles in various applications, techniques like Self-Refine Prompting will be crucial in realizing their full potential while addressing current limitations. This work contributes to the growing body of research on LLM optimization and self-improvement, offering both practical applications for current models and insights to guide future developments in artificial intelligence and natural language processing.

REFERENCES

1. Brown TB, Mann B, Ryder N, Subbiah M, Kaplan J, Dhariwal P, et al. Language models are few-shot learners. arXiv preprint arXiv:2005.14165. 2020 May 28.
2. Chowdhery A, Narang S, Devlin J, Bosma M, Mishra G, Roberts A, et al. PaLM: Scaling language modeling with Pathways. arXiv preprint arXiv:2204.02311. 2022 Apr 5.
3. Anthropic. Claude: AI assistant [Internet]. San Francisco: Anthropic; 2023 [cited 2024 Jul 31]. Available from: https://www.anthropic.com
4. Dou ZY, Forbes M, Koncel-Kedziorski R, Smith NA, Neubig G. Scarecrow: A framework for scrutinizing machine text. arXiv preprint arXiv:2107.01294. 2021 Jul 2.
5. Wang A, Pruksachatkun Y, Nangia N, Singh A, Michael J, Hill F, et al. SuperGLUE: A stickier benchmark for general-purpose language understanding systems. arXiv preprint arXiv:1905.00537. 2019 May 1.
6. Howard J, Ruder S. Universal language model fine-tuning for text classification. arXiv preprint arXiv:1801.06146. 2018 Jan 18.
7. Brown TB, Mann B, Ryder N, Subbiah M, Kaplan J, Dhariwal P, et al. Language models are few-shot learners. arXiv preprint arXiv:2005.14165. 2020 May 28.
8. Liu P, Yuan W, Fu J, Jiang Z, Hayashi H, Neubig G. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. arXiv preprint arXiv:2107.13586. 2021 Jul 28.
9. Jurafsky D, Martin JH. Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition. Upper Saddle River (NJ): Prentice Hall; 2009.
10. Rabiner LR. A tutorial on hidden Markov models and selected applications in speech recognition. Proceedings of the IEEE. 1989 Feb;77(2):257-86.
11. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is all you need. In: Advances in neural information processing systems; 2017 Dec. p. 5998-6008.
12. Devlin J, Chang MW, Lee K, Toutanova K. BERT: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805. 2018 Oct 11.
13. Radford A, Narasimhan K, Salimans T, Sutskever I. Improving language understanding by generative pre-training [Internet]. OpenAI; 2018 [cited 2024 Jul 31]. Available from:

https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf
14. Kaplan J, McCandlish S, Henighan T, Brown TB, Chess B, Child R, et al. Scaling laws for neural language models. arXiv preprint arXiv:2001.08361. 2020 Jan 23.
15. Zhang Y, Sun S, Galley M, Chen YC, Brockett C, Gao X, et al. DialoGPT: Large-scale generative pre-training for conversational response generation. arXiv preprint arXiv:1911.00536. 2019 Nov 1.
16. Bosselut A, Rashkin H, Sap M, Malaviya C, Celikyilmaz A, Choi Y. COMET: Commonsense transformers for automatic knowledge graph construction. arXiv preprint arXiv:1906.05317. 2019 Jun 12.
17. Yu D, Sun K, Cardie C, Yu D. Dialogue-based relation extraction. arXiv preprint arXiv:2004.08056. 2020 Apr 17.
18. Ziegler DM, Stiennon N, Wu J, Brown TB, Radford A, Amodei D, et al. Fine-tuning language models from human preferences. arXiv preprint arXiv:1909.08593. 2019 Sep 18.
19. Desai S, Durrett G. Calibration of pre-trained transformers. arXiv preprint arXiv:2003.07892. 2020 Mar 17.
20. Guo C, Pleiss G, Sun Y, Weinberger KQ. On calibration of modern neural networks. In: International Conference on Machine Learning; 2017 Aug. p. 1321-30.
21. Jiang Z, Xu FF, Araki J, Neubig G. How can we know what language models know? Transactions of the Association for Computational Linguistics. 2020 Dec;8:423-38.
22. Hermann KM, Kocisky T, Grefenstette E, Espeholt L, Kay W, Suleyman M, et al. Teaching machines to read and comprehend. In: Advances in neural information processing systems; 2015 Dec.
23. Rajpurkar P, Zhang J, Lopyrev K, Liang P. SQuAD: 100,000+ questions for machine comprehension of text. arXiv preprint arXiv:1606.05250. 2016 Jun 16.
24. Fan A, Lewis M, Dauphin Y. Hierarchical neural story generation. arXiv preprint arXiv:1805.04833. 2018 May 13.
25. Lin CY. Rouge: A package for automatic evaluation of summaries. In: Text summarization branches out; 2004 Jul. p. 74-81.
26. Jelinek F, Mercer RL, Bahl LR, Baker JK. Perplexity—a measure of the difficulty of speech recognition tasks. The Journal of the Acoustical Society of America. 1977 Nov;62(S1):S63.
27. Likert R. A technique for the measurement of attitudes. Archives of psychology. 1932 Jun.
28. Tsimpoukelli M, Menick J, Cabi S, Eslami SA, Vinyals O, Hill F. Multimodal few-shot learning with frozen language models. In: Advances in Neural Information Processing Systems; 2021 Dec. p. 200-12.
29. Liu W, Zhou P, Wang Z, Zhao Z, Deng H, Ju Q. K-BERT: Enabling language representation with knowledge graph. In: Proceedings of the AAAI Conference on Artificial Intelligence; 2020 Apr. p. 2901-8.
30. Petroni F, Rocktäschel T, Riedel S, Lewis P, Bakhtin A, Wu Y, et al. Language models as knowledge bases? arXiv preprint arXiv:1909.01066. 2019 Sep 3.
31. Zhang Y, Sun S, Galley M, Chen YC, Brockett C, Gao X, et al. DialoGPT: Large-scale generative pre-training for conversational response generation. arXiv preprint arXiv:1911.00536. 2019 Nov 1.
32. Liu Y, Gu J, Goyal N, Li X, Edunov S, Ghazvininejad M, et al. Multilingual denoising pre-training for neural machine translation. Transactions of the Association for Computational Linguistics. 2020 Dec;8:726-42.

International Journal of Core Engineering & Management
Volume-7, Issue-10, 2024 ISSN No: 2348-9510

33. Bommasani R, Hudson DA, Adeli E, Altman R, Arora S, von Arx S, et al. On the opportunities and risks of foundation models. arXiv preprint arXiv:2108.07258. 2021 Aug 16.
34. Solaiman I, Dennison C. Process for adapting language models to society (PALMS) with values-targeted datasets. arXiv preprint arXiv:2106.10328. 2021 Jun 18.
35. Madaan A, Tandon N, Gupta P, Hallinan S, Gao L, Wiegreffe S, et al. Self-Refine: Iterative refinement with self-feedback. Advances in Neural Information Processing Systems. 2023june;36:46534-94.
36. Bhaskarjit. Reflexion: Language Agents with Verbal Reinforcement Learning [Internet]. Substack.com. Visual GenAI Summary; 2023Aug. Available from: https://visualsummary.substack.com/p/reflexion-language-agents-with-verbal