

ABSTRACT

Indonesian palm oil regulatory documents feature lengthy texts densely packed with legal terminology, making it difficult for stakeholders to efficiently grasp their substance. This study explores the effectiveness of prompt engineering strategies for abstractive summarization of palm oil regulations using two Large Language Models (LLMs), GPT-5.4-mini and Llama 3.1:8B, tested on 157 regulatory documents. The experiment combined three prompting strategies (zero-shot, one-shot, few-shot) with four context configuration variants derived from two components, persona and domain knowledge, ranging from a no-context baseline (*general*) to their individual and combined forms, producing 24 scenarios and 3,768 summaries evaluated using BERTScore, BLEU, ROUGE-L, and METEOR. Results show that zero-shot dominates three out of four metrics (BLEU 0.1469; ROUGE-L 0.3267; BERTScore 0.7571), as adding examples in one-shot and few-shot settings risks overprompting toward the sample documents' specific writing style. The persona configuration consistently yields the best results across all metrics, while domain knowledge variants degrade performance due to excessive context load. GPT-5.4-mini outperforms Llama 3.1:8B across all metrics. Overall, the zero-shot persona combination with GPT-5.4-mini proves to be the optimal configuration, achieving the highest BERTScore of 0.7847, BLEU 0.1786, and ROUGE-L 0.3675.

Keywords: Abstractive Summarization, *Prompt* Engineering, Large Language Model.