

# Context-Aware Retrieval and RAG Pipelines in Agentic Code Completion Trade-offs

Assignee Research

June 2, 2026

## Abstract

This report synthesises findings from 3 peer-reviewed papers addressing the following research question: What is the trade-off between token generation speed and API call precision in agentic code completion systems using context-aware retrieval versus standard RAG pipelines. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.2/10. This report is a machine-generated literature synthesis and does not constitute original research.

## 1 Introduction

This paper examines: A Comprehensive Overview of Large Language Models. Research question: What is the trade-off between token generation speed and API call precision in agentic code completion systems using context-aware retrieval versus standard RAG pipelines?.

## 2 Methodology

Systematic literature search across multiple databases yielded 3 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.2/10.

## 3 Results

3 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 3.2/10.

## 4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

## References

- <https://doi.org/10.1039/d4sc03921a>
- <https://doi.org/10.48550/arxiv.2307.06435>
- <https://doi.org/10.48550/arxiv.2307.10169>