

Vendi-RAG Performance Across Domains: Adaptive Diversity-Weight Tuning in Code and Multimodal Tasks

Assignee Research

May 30, 2026

Abstract

This report synthesises findings from 11 peer-reviewed papers addressing the following research question: How does the performance of Vendi-RAG with adaptive diversity-weight tuning vary across different domains (e.g., code generation with HumanEval vs. multimodal reasoning with MMQA) when measured by. Understanding and reasoning over diagrams is a fundamental aspect of human intelligence. While Large Multimodal Models (LMMs) have demonstrated impressive capabilities across various tasks, existing benchmarks lack comprehensive evaluation of their diagram interpretation and. 12 claims were extracted from source literature; 1 was independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: HumanEval-V: Benchmarking High-Level Visual Reasoning with Complex Diagrams in Coding Tasks. Research question: How does the performance of Vendi-RAG with adaptive diversity-weight tuning vary across different domains (e.g., code generation with HumanEval vs. multimodal reasoning with MMQA) when measured by both accuracy and inference efficiency?.

2 Methodology

Systematic literature search across multiple databases yielded 11 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.5/10.

3 Results

11 papers retrieved. 12 claims extracted; 1 independently verified. Quality review score: 4.5/10.

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.

5 Extracted Claims

Claim	Verified	Confidence
HumanEval-V consists of 253 human-annotated coding tasks.	✓	0.16
Each task in HumanEval-V features a diagram encoding the problem context, a function signature defining the task’s input	×	0.07
The top-performing model, Claude 3.5 Sonnet, achieves 36.8% pass@1 on HumanEval-V.	×	0.10
The best open-weight model, Pixtral 124B, reaches 21.3% pass@1 on HumanEval-V.	×	0.03
Claude 3.5 Sonnet achieves a 74.3% pass rate with 100 samples on HumanEval-V.	×	0.05
Claude 3.5 Sonnet can reach 55.3% pass@1 with four self-refining iterations based on test case execution feedback on Hum	×	0.04
HumanEval-V offers a more diverse and complex set of diagrams spanning six task types, demanding versatile capabilities	×	0.14
HumanEval-V uses code generation tasks for evaluation instead of the multiple-choice or short-answer questions commonly	×	0.11
The visual context must be essential for solving the task in HumanEval-V, with all relevant information contained in a s	×	0.05
Tasks in HumanEval-V should be designed around the visual context with minimal textual description.	×	0.04
HumanEval-V utilizes a two-stage evaluation pipeline that supports LMMs with limited coding abilities by first prompting	×	0.07
Through extensive experiments with 22 LMMs, key findings were observed on HumanEval-V.	×	0.14

References

- <http://arxiv.org/abs/2102.07660v2>
- <http://arxiv.org/abs/2603.03437v1>
- <http://arxiv.org/abs/2410.12381v3>