

LLaVA-v1.6-Mistral-7B Benchmark Performance Across Reasoning Mathematics Coding and Language Tasks

Assignee Research

June 6, 2026

Abstract

This report synthesises findings from 15 peer-reviewed papers addressing the following research question: What are the benchmark performance scores of llava-v1.6-mistral-7b on reasoning mathematics coding and language understanding tasks. 17 claims were extracted from source literature; 2 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.8/10. 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: What are the benchmark performance scores of llava-v1.6-mistral-7b on reasoning mathematics coding and language understanding tasks.

2 Methodology

Systematic literature search across multiple databases yielded 15 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.8/10.

3 Results

15 papers retrieved. 17 claims extracted; 2 independently verified. Quality review score: 4.8/10.

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.	×	0.11
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.	×	0.10
HumanEval-V demands versatile capabilities for diagram understanding and reasoning.	×	0.09
HumanEval-V uses code generation tasks for evaluation instead of multiple-choice or short-answer questions.	×	0.13
Current LMMs exhibit stronger vision-to-language alignment than vision-to-code.	×	0.07
Current LMMs still have difficulty understanding diagrams that are trivial for humans, particularly understanding spatia	✓	0.21
Each coding task in HumanEval-V includes a single diagram D providing visual context, a Python function signature σ with	×	0.07
The visual context must be essential for solving the task, with all relevant information contained in a single image.	×	0.05
Tasks should be designed around the visual context with minimal textual description.	×	0.04
Test cases could rigorously verify whether the model captures all critical visual information, rather than relying on si	×	0.03
The evaluation pipeline supports LMMs with limited coding abilities by first prompting them to generate a structured dia	×	0.07
Extensive experiments with 22 LMMs were conducted.	×	0.15

References

- <http://arxiv.org/abs/2509.25160v1>
- <http://arxiv.org/abs/2410.12381v3>
- <http://arxiv.org/abs/2210.09261v1>