



ScratchEval: Are GPT-4o Smarter than My Child? Evaluating Large Multimodal Models with Visual Programming Challenges

Rao Fu[♡], Ziyang Luo[♡], Hongzhan Lin[♡], Zhen Ye[♣], Jing Ma^{♡*}

[♡]Hong Kong Baptist University, [♣]Hong Kong University of Science and Technology
maging@comp.hkbu.edu.hk

Abstract

Recent advancements in large multimodal models (LMMs) have showcased impressive code generation capabilities, primarily evaluated through image-to-code benchmarks. However, these benchmarks are limited to specific visual programming scenarios where the logic reasoning and the multimodal understanding capacities are split apart. To fill this gap, we propose **ScratchEval**, a novel benchmark designed to evaluate the visual programming reasoning ability of LMMs. **ScratchEval** is based on Scratch, a block-based visual programming language widely used in children’s programming education. By integrating visual elements and embedded programming logic, **ScratchEval** requires the model to process both visual information and code structure, thereby comprehensively evaluating its programming intent understanding ability. Our evaluation approach goes beyond the traditional image-to-code mapping and focuses on unified logical thinking and problem-solving abilities, providing a more comprehensive and challenging framework for evaluating the visual programming ability of LMMs. **ScratchEval** not only fills the gap in existing evaluation methods, but also provides new insights for the future development of LMMs in the field of visual programming. Our benchmark can be accessed at <https://github.com/HKBUNLP/ScratchEval>.

1 Introduction

Recently, Large Multimodal Models (LMMs) such as GPT-4o (OpenAI, 2023), Gemini (Anil et al., 2023), and Claude (Anthropic, 2023) have shown remarkable capabilities in multimodal understanding (Chen et al., 2024a; Lin et al., 2024; Wang et al., 2024b; Luo et al., 2024; Yu et al., 2024). To assess their abilities, several comprehensive benchmarks have been introduced, including MMMU (Yue

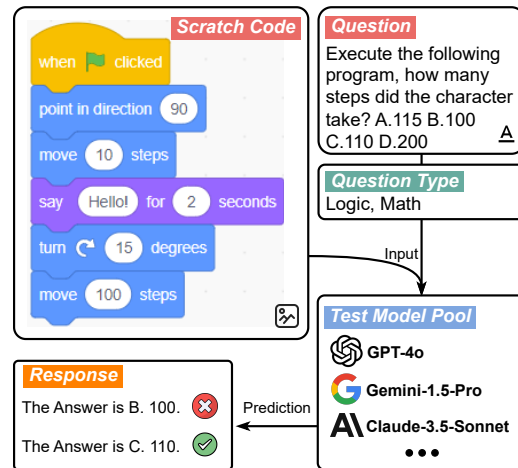


Figure 1: The illustration of the evaluation process for **ScratchEval**.

et al., 2023), MME (Fu et al., 2023), MathVista (Lu et al., 2024), and MMBench (Liu et al., 2023). These benchmarks primarily focus on evaluating core multimodal skills of LMMs, such as object detection, OCR, and visual reasoning. The evaluations provide deeper insights into the strengths and limitations of LMMs.

In addition to general multimodal understanding tasks, recent works such as MMCode (Li et al., 2024), Design2Code (Si et al., 2024), Plot2Code (Wu et al., 2024), and CharMimic (Shi et al., 2024) focus on assessing the visual programming reasoning abilities of LMMs. Most of the previous work focuses on specific scenarios, such as converting matplotlib images to Python code, generating code based on diagrams of algorithmic problems, or even generating HTML code from web page screenshots. Although these studies include visual elements, the diversity of input is relatively limited, mainly focusing on a single mapping from image to code, but ignoring cases where the programming logic is inherent in images.

*Corresponding Authors.

In this paper, we argue that it is imperative to evaluate the visual programming capacity of LMMs by unifying visual understanding and logical reasoning. Inspired by children’s coding education (Pérez-Marín et al., 2020), using a graphical programming way, allows the assessment to focus more on logical thinking and problem-solving skills, rather than traditional programming languages that may be plagued by syntax errors. Thus we aim to combine visual elements with programming logic, requiring LMMs to process both visual information and code structure.

To this end, we introduce **ScratchEval** as illustrated in Figure 1, a novel benchmark designed to assess LMMs’ visual programming reasoning abilities by integrating visual elements with embedded programming logic. ScratchEval is based on Scratch (Dasgupta and Hill, 2017), a popular block-based visual programming language widely used as an educational tool for children aged 8 to 16. It allows users to create projects through a drag-and-drop block interface, offering a visual approach to coding. By leveraging graphical code, our evaluation focuses on the complexity of multimodal input, where the model must understand the image, graphical programming language, and underlying logic, showcasing a comprehensive grasp of programming intent.

On ScratchEval, we tested multiple existing open-source and closed-source LMMs and studied the impact of different prompting strategies on model performance. Finally, we conducted a case study to analyze the performance bottleneck of the model. Through our research, we found that the existing state-of-the-art LMMs still fail to achieve high performance on our proposed benchmark, which shows the inadequacy of existing models in visual code reasoning capabilities and also points out the direction for further research.

2 ScratchEval

All our data is manually collected and cleaned by experts from public question banks on the web. We organized the data into 305 multiple-choice questions, each with a problem description, options, and a picture containing the Scratch script and other necessary information.

Our test benchmark consists of two components: Chinese and English data. Both sections are identical in quantity and content, but the questions and Scratch script images are in their respective lan-

Task	Number
Math	133
Logical thinking	99
Graphic perception	59
Spacial perception	43
All	305

Table 1: Data volume of the four tasks, each question examines at most two types of the tasks.

guages. This approach evaluates the visual reasoning capabilities of various models across different linguistic contexts, allowing us to assess how language-specific factors influence performance in interpreting visual information in Scratch programming. By comparing results from both datasets, we gain insights into the models’ cross-linguistic robustness and adaptability.

2.1 Data analysis

Based on the content of the questions, we categorized them into four domains: mathematics, logical thinking, graphic perception, and spatial perception. The specific distribution of questions across these categories is presented in Table 1. It is important to note that some questions evaluate multiple abilities, and therefore, each question is assigned to at most two categories. The characteristics of each category are as follows:

Mathematics tasks encompass simple arithmetic problems typically encountered in elementary and junior high school curricula. These tasks assess the model’s ability to solve basic mathematical problems.

Logical thinking tasks evaluate the model’s capacity for logical reasoning based on provided Scratch scripts. These scripts are designed for children and are generally comprehensible even to those unfamiliar with the Scratch programming environment.

Graphic perception tasks examine the model’s understanding of graphics. These may involve selecting graphics that correspond to a given script or inferring the output of a simple drawing program.

Spatial perception tasks assess the model’s ability to determine the final position and orientation of a character based on a movement program.

This categorization enables thorough assessment of models’ visual code reasoning abilities across cognitive domains.

2.2 Evaluation Methodology

The evaluation process consists of three stages: 1) generating answers, 2) extracting answers, and 3) calculating scores.

First, the tested LMM generates answers based on the input query, which includes questions, options, image data, and a system prompt. After our experiments, the system prompt we set can help us greatly simplify the output of the model. Finally, the extracted answers are normalized to the required answer format option letters, and the target metric score is calculated. Using the fact that the examples in ScratchEval are multiple-choice questions with text answers, the accuracy score is used as a metric for deterministic evaluation.

3 Experiments

3.1 Experiment setup

We evaluate a total of 10 LMMs on ScratchEval under two setups: (a) Closed-source LMMs, including Gemini-1.5-Pro (Reid et al., 2024), GPT-4-Turbo (Achiam et al., 2023), GPT-4o, and Claude-3.5-Sonnet; (b) Open-source LMMs, including Qwen2-VL (Wang et al., 2024a), LLaVA-v1.6 (Liu et al., 2024), InternVL2 (Chen et al., 2024b), Pixtral (Agrawal et al., 2024), MiniCPM-v2.6 (Yao et al., 2024) and Molmo (Deitke et al., 2024). We use the accuracy as the evaluation metric. We provide implementation details in the Appendix §A.1.

3.2 Experiment analysis

We evaluated the performance of 10 state-of-the-art LMMs by drawing the practice of the LMSYS Chatbot Arena leaderboard on our proposed ScratchEval benchmark, incorporating both Chinese and English data. The experimental results on English data are presented in Table 2. To conduct a detailed analysis of the LMMs’ capabilities, we categorized the questions into four domains: mathematics, logical thinking, graphic perception, and spatial perception.

The results reveal significant performance variations across models in each category, with most models surpassing the 25% random guessing threshold. This indicates that LMMs possess some visual code reasoning capabilities, enabling them to process visual information alongside language comprehension.

Gemini-1.5-Pro demonstrated superior performance, achieving the highest scores across all categories. However, most other models struggled to

exceed 50% accuracy, highlighting current limitations in LMMs regarding visual code reasoning. We attribute this to a lack of high-quality visual-language paired data during training, as larger models like Gemini-1.5-pro and GPT-4o performed better. Additionally, the model’s vision tokenizer may influence its visual reasoning capabilities.

Most models underperformed in mathematical and logical reasoning tasks, suggesting a deficiency in multi-step reasoning. Conversely, LMMs exhibited better performance in graphic and spatial perception tasks, demonstrating an understanding of concepts such as orientation and distance, which they can leverage for reasoning to some extent. The experimental results on Chinese data can be found in the Appendix §A.5.

3.3 Prompting strategies study

We investigated the impact of prompt engineering on the visual code reasoning capabilities of models using our test benchmark. Previous studies, such as CoT (Wei et al., 2023), have shown that appropriate prompting can enhance the performance of large language models. However, its effectiveness for multimodal large language models remains underexplored. To address this, we selected four models and applied three prompting strategies to examine their influence on reasoning abilities.

The prompting strategies employed were: (1) Original prompt ("no-CoT"): using raw data as prompts. (2) zero-shot CoT ("CoT"): Chain of Thought prompting, appending "Let’s think step by step." to each question for more comprehensive analysis. (3) eCoT: Inspired by (Ghosal et al., 2024), we implemented eCoT, which requires a detailed examination during the CoT process by appending "Let’s explain the picture and think step by step." to each question.

We found that CoT and eCoT techniques significantly enhanced the models’ visual code reasoning capabilities, with CoT prompting improving performance by 10% to 20%. However, no model achieved overall accuracy exceeding 70%, indicating substantial room for improvement. Additionally, eCoT yielded relatively minor improvements compared to CoT, suggesting that describing the image may hinder the model’s visual code reasoning capabilities. Detailed experimental data can be found in the Appendix §A.5

Models	Size	All	Math	Logical Thinking	Graphic Perception	Spatial Perception
<i>Proprietary Models</i>						
Gemini-1.5-Pro	-	52.8	55.3	49.5	47.5	59.5
GPT-4o	-	43.9	44.7	42.4	45.8	50.0
GPT-4-Turbo	-	40.7	39.4	44.4	37.3	43.0
Claude-3.5-Sonnet	-	40.3	45.5	37.3	35.6	35.7
<i>Open-Source Models</i>						
Qwen2-VL	72B	45.0	50.0	42.4	45.8	40.5
LLaVA-v1.6	34B	26.5	21.2	30.3	35.6	26.2
InternVL2	26B	22.3	25.6	18.2	20.3	21.4
Pixtral	12B	34.1	34.1	34.3	32.2	28.6
MiniCPM-v2.6	8B	30.0	28.0	31.3	39.0	31.0
Molmo	7B	31.2	32.6	29.3	33.9	26.2

Table 2: Accuracy (%) of ten state-of-the-art LMMs on the English data of ScratchEval benchmark, tested across multiple cognitive abilities: math, logical thinking, graphic perception, and spatial perception.

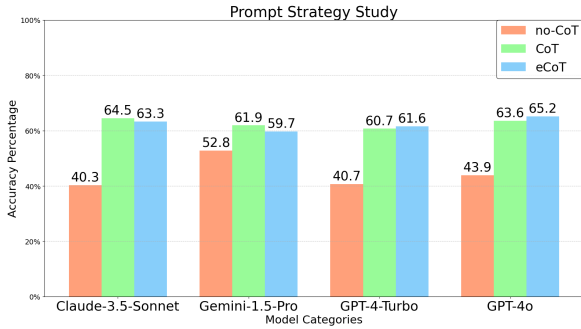


Figure 2: Models’s performance under different prompting strategies.

3.4 Case study

To better understand the model’s behavior, we selected several examples where Gemini-1.5-Pro made mistakes for a case study. Overall, Gemini-1.5-Pro is the best-performing model in ScratchEval. By studying its behavior, we aim to explain why ScratchEval is challenging for most models.

We chose representative examples for Gemini-1.5-Pro’s case study, as shown in Figure 3. We specifically selected examples that failed across all three prompting strategies mentioned earlier, allowing us to observe Gemini-1.5-Pro’s deficiencies in certain areas.

As shown in Figure 3, Gemini-1.5-Pro with CoT accurately identified image content but hallucinated during reasoning. With eCoT, it described the image but misinterpreted symbols, leading to incorrect inferences.

These cases reveal that while Gemini-1.5-Pro excels in reasoning and basic math/logic problems,

Scratch Code

Question

What are the values of the two blocks in the picture below?
A.false true B.false false C.true true D.true false
Correct Answer: A

eCoT Answer

The script evaluates in two parts:Part 1: Left Side* $50 = 50$: This comparison is true (50 does equal 50).* $50 = 50$: This comparison is also true.* true or true*: Since at least one side of the "or" is true, the entire left side evaluates to true.Part 2: Right Side* $10 \wedge$ of 2: This is using the "letter of" block, which doesn't have a standard mathematical meaning. It's likely meant to be a custom function or a misunderstanding of the block's purpose. Therefore, we cannot determine a definite true/false value for this part.* $= 50$: Regardless of the result of the previous calculation, comparing it to 50 will either be true or false.OverallSince the right side's outcome is uncertain, we can only be sure of the left side's value, which is true.Therefore, none of the provided options (A, B, C, D) are definitively correct.

CoT Answer

Let's break down the script: $50 = 50$: This is true. $50 = 50$: This is also true.or: This means if either side is true, the whole thing is true. Since both sides are true, this part is true.* $10 \wedge$ of 2: This is calculating 10 to the power of 2, which is $100 \cdot 100 = 50$: This is false.* and: This means both sides must be true for the whole thing to be true.Putting it together:The first part ($50 = 50$ or $50 = 50$) is true.The second part ($10 \wedge$ of 2 = 50) is false.Since the two parts are connected by "and", and the second part is false, the entire statement is false.Therefore, the values of the two blocks are:* true (for the first part)* false (for the second part)The correct answer is D. true false

Figure 3: A Gemini-1.5-Pro mistake case. The error areas are marked in red.

it struggles with subtle image distinctions. Visual encoders and hallucinations remain the main bottlenecks that restrict the model’s reasoning capabilities. Additional examples are provided in Appendix A.5.

4 Conclusion

In this work, we present ScratchEval, a benchmark that uses the Scratch language to systematically evaluate the visual programming capabilities of state-of-the-art LMMs. Our evaluation of 10 representative LMMs indicates that while these models show some visual comprehension, they struggle

with visual code reasoning. This highlights the need for research on models that integrate visual perception with logical thinking. ScratchEval provides a foundation for future studies aimed at enhancing AI systems' visual reasoning capabilities, bridging the gap between visual understanding and logical reasoning in LMMs.

5 Limitations

Although our proposed ScratchEval helps us to evaluate the visual reasoning ability of existing LMMs, we recognize that our work still has several important limitations: (1) Due to the difficulty of LMMs to directly operate graphical programming languages, in order to use graphical programming to examine the model's visual programming abilities, we model the problem as Multiple choice questions. (2) the narrow domain focus of our benchmark, concentrating solely on visual programming abilities, limits the generalizability of our findings. The results obtained cannot be extrapolated to assess other competencies of LMMs. These limitations underscore the need for continued research and development of more comprehensive evaluation methodologies for large multimodal models.

6 Acknowledge

This work is partially supported by National Natural Science Foundation of China Young Scientists Fund(No. 62206233) and Hong Kong RGC ECS (No. 22200722).

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Pravesh Agrawal, Szymon Antoniak, Emma Bou Hanna, Devendra Chaplot, Jessica Chudnovsky, Saurabh Garg, Theophile Gervet, Soham Ghosh, Amal'lie HÅ'liou, Paul Jacob, Albert Q. Jiang, Timoth-Å'e Lacroix, Guillaume Lample, Diego Las Casas, Thibaut Lavril, Teven Le Scao, Andy Lo, William Marshall, Louis Martin, Arthur Mensch, Pavankumar Muddireddy, Valera Nemychnikova, Marie Pellat, Patrick Von Platen, Nikhil Raghuraman, Baptiste RoziÅ're, Alexandre Sablayrolles, Lucile Saulnier, Romain Sauvestre, Wendy Shang, Roman Soletskyi, Lawrence Stewart, Pierre Stock, Joachim Studnia, Sandeep Subramanian, Sagar Vaze, and Thomas Wang. 2024. *Pixtral 12b*. *Preprint*, arXiv:2410.07073.
- Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M. Dai, Anja Hauth, Katie Millican, David Silver, Slav Petrov, Melvin Johnson, Ioannis Antonoglou, Julian Schrittwieser, Amelia Glaese, Jilin Chen, Emily Pitler, Timothy P. Lillcrap, Angeliki Lazaridou, Orhan Firat, James Molloy, Michael Isard, Paul Ronald Barham, Tom Hennigan, Benjamin Lee, Fabio Viola, Malcolm Reynolds, Yuanzhong Xu, Ryan Doherty, Eli Collins, Clemens Meyer, Eliza Rutherford, Erica Moreira, Kareem Ayoub, Megha Goel, George Tucker, Enrique Pi-queras, Maxim Krikun, Iain Barr, Nikolay Savinov, Ivo Danihelka, Becca Roelofs, Anaïs White, Anders Andreassen, Tamara von Glehn, Lakshman Yagati, Mehran Kazemi, Lucas Gonzalez, Misha Khalman, Jakob Sygnowski, and et al. 2023. *Gemini: A family of highly capable multimodal models*. *CoRR*, abs/2312.11805.
- Anthropic. 2023. Claude: A family of large language models. <https://www.anthropic.com/claude>.
- Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi Wang, Yu Qiao, Dahua Lin, and Feng Zhao. 2024a. *Are we on the right way for evaluating large vision-language models?* *CoRR*, abs/2403.20330.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. 2024b. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24185–24198.
- Sayamindu Dasgupta and Benjamin Mako Hill. 2017. Scratch community blocks: Supporting children as data scientists. In *Proceedings of the 2017 CHI conference on human factors in computing systems*, pages 3620–3631.
- Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, Jiasen Lu, Taira Anderson, Erin Bransom, Kiana Ehsani, Huong Ngo, YenSung Chen, Ajay Patel, Mark Yatskar, Chris Callison-Burch, Andrew Head, Rose Hendrix, Favyen Bastani, Eli Vander-Bilt, Nathan Lambert, Yvonne Chou, Arnavi Chheda, Jenna Sparks, Sam Skjonsberg, Michael Schmitz, Aaron Sarnat, Byron Bischoff, Pete Walsh, Chris Newell, Piper Wolters, Tanmay Gupta, Kuo-Hao Zeng, Jon Borchardt, Dirk Groeneveld, Jen Dumas, Crystal Nam, Sophie Lebrecht, Caitlin Wiltliff, Carissa Schoenick, Oscar Michel, Ranjay Krishna, Luca Weihs, Noah A. Smith, Hannaneh Hajishirzi, Ross Girshick, Ali Farhadi, and Aniruddha Kembhavi. 2024. *Molmo and pixmo: Open weights and open data for state-of-the-art multimodal models*. *Preprint*, arXiv:2409.17146.
- Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Zhenyu Qiu, Wei Lin, Jin-

- rui Yang, Xiawu Zheng, Ke Li, Xing Sun, and Rongrong Ji. 2023. [MME: A comprehensive evaluation benchmark for multimodal large language models](#). *CoRR*, abs/2306.13394.
- Deepanway Ghosal, Vernon Toh Yan Han, Chia Yew Ken, and Soujanya Poria. 2024. [Are language models puzzle prodigies? algorithmic puzzles unveil serious challenges in multimodal reasoning](#). *Preprint*, arXiv:2403.03864.
- Kaixin Li, Yuchen Tian, Qisheng Hu, Ziyang Luo, and Jing Ma. 2024. [Mmcode: Evaluating multi-modal code large language models with visually rich programming problems](#). *CoRR*, abs/2404.09486.
- Hongzhan Lin, Ziyang Luo, Bo Wang, Ruichao Yang, and Jing Ma. 2024. [Goat-bench: Safety insights to large multimodal models through meme-based social abuse](#). *Preprint*, arXiv:2401.01523.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024. [Llava-next: Improved reasoning, ocr, and world knowledge](#).
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, Kai Chen, and Dahua Lin. 2023. [Mmbench: Is your multi-modal model an all-around player?](#) *CoRR*, abs/2307.06281.
- Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. 2024. [Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Ziyang Luo, Haoning Wu, Dongxu Li, Jing Ma, Mohan Kankanhalli, and Junnan Li. 2024. [Videoautoarena: An automated arena for evaluating large multimodal models in video analysis through user simulation](#). *Preprint*, arXiv:2411.13281.
- OpenAI. 2023. [GPT-4 technical report](#). *CoRR*, abs/2303.08774.
- Diana Pérez-Marín, Raquel Hijón-Neira, Adrián Babelo, and Celeste Pizarro. 2020. [Can computational thinking be improved by using a methodology based on metaphors and scratch to teach computer programming to children?](#) *Computers in Human Behavior*, 105:105849.
- Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy Lillicrap, Jean-baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, et al. 2024. [Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context](#). *arXiv preprint arXiv:2403.05530*.
- Chufan Shi, Cheng Yang, Yaxin Liu, Bo Shui, Junjie Wang, Mohan Jing, Linran Xu, Xinyu Zhu, Siheng Li, Yuxiang Zhang, Gongye Liu, Xiaomei Nie, Deng Cai, and Yujiu Yang. 2024. [Chartmimic: Evaluating Imm’s cross-modal reasoning capability via chart-to-code generation](#). *CoRR*, abs/2406.09961.
- Chenglei Si, Yanzhe Zhang, Zhengyuan Yang, Ruibo Liu, and Diyi Yang. 2024. [Design2code: How far are we from automating front-end engineering?](#) *CoRR*, abs/2403.03163.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. 2024a. [Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution](#). *Preprint*, arXiv:2409.12191.
- Shengkang Wang, Hongzhan Lin, Ziyang Luo, Zhen Ye, Guang Chen, and Jing Ma. 2024b. [Mfc-bench: Benchmarking multimodal fact-checking with large vision-language models](#). *Preprint*, arXiv:2406.11288.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. [Chain-of-thought prompting elicits reasoning in large language models](#). *Preprint*, arXiv:2201.11903.
- Chengyue Wu, Yixiao Ge, Qiushan Guo, Jiahao Wang, Zhixuan Liang, Zeyu Lu, Ying Shan, and Ping Luo. 2024. [Plot2code: A comprehensive benchmark for evaluating multi-modal large language models in code generation from scientific plots](#). *CoRR*, abs/2405.07990.
- Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li, Weilin Zhao, Zhihui He, et al. 2024. [Minicpm-v: A gpt-4v level mllm on your phone](#). *arXiv preprint arXiv:2408.01800*.
- Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. 2024. [Mm-vet: Evaluating large multimodal models for integrated capabilities](#). In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net.
- Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, Cong Wei, Botao Yu, Ruibin Yuan, Renliang Sun, Ming Yin, Boyuan Zheng, Zhenzhu Yang, Yibo Liu, Wenhao Huang, Huan Sun, Yu Su, and Wenhui Chen. 2023. [MMMU: A massive multi-discipline multimodal understanding and reasoning benchmark for expert AGI](#). *CoRR*, abs/2311.16502.

A Appendix

A.1 Experiments setup

In our study, we conducted comprehensive evaluations of 10 state-of-the-art Large Multimodal Models (LMMs) on the ScratchEval benchmark. The following models were included in our experiments:

- Gemini-1.5-Pro-Exp-0827
- GPT-4o-2024-05-13
- Claude 3.5 Sonnet
- GPT-4-Turbo-2024-04-09
- Qwen2-VL-72b-Instruct
- InternVL2-26b
- LLaVA-v1.6-34B
- MiniCPM-V 2_6
- Pixtral-12b-2409
- Molmo-7B-D-0924

All models were evaluated using their respective latest versions available at the time of the experiment. To ensure consistency and reproducibility across all tests, we maintained a constant temperature setting of 0 for all models. This setting was chosen to produce deterministic outputs and facilitate direct comparisons between models.

For each model, depending on the task being performed, we use specific system prompts to explain the next task to the model. These system prompts are as follows:

- For no-CoT tasks: "According to the displayed Scratch script and the given question, please choose a correct answer from the four options ABCD. You only need to find the correct option, and no analysis is required. "
- For CoT tasks: "According to the displayed Scratch script and the given question, please choose a correct answer from the four options ABCD. "
- For eCoT tasks: "According to the displayed Scratch script and the given question, please choose a correct answer from the four options ABCD. "

The system prompts when executing Chinese tasks are the translations of the above corresponding tasks.

A.2 Chinese data experiments

In Table 3, We can see that the performance of most models is basically the same as in the English task, while some models perform better. We believe this is because some models use more Chinese data during training.

A.3 Data example

In Figure 6, Figure 7, Figure 8 and Figure 9, we show data for mathematics, logical thinking, graphic perception, and Spatial perception as examples. Each example includes the corresponding Chinese and English scripts, questions, and correct answers.

A.4 Prompt strategic study data

In Figure 5, we provide more data on the model performance under different prompt strategies, which are also consistent with the views we put forward in the main text.

A.5 Examples in case study

In Figure 4, We show two cases where Gemini-1.5-Pro makes mistakes, and these two cases also illustrate the conclusions we stated in the main text.

A.6 Potential Risks

While our benchmark for LMMs, which evaluates models using Scratch visual programming questions, poses no direct risks, potential concerns include the possibility of models overfitting to specific visual programming patterns, reducing their generalization capabilities. Additionally, the reliance on Scratch could limit the applicability of results to broader real-world tasks that use different programming interfaces.

A.7 Creators Of Artifacts

The source data for our benchmark is derived from the China Lanqiao Cup National Software and Information Technology Professional Talent Competition <https://www.lanqiaoqingshao.cn/home> (Chinese website). To adapt this data for our benchmark, we enlisted the help of domain experts to reannotate and refine the original dataset, ensuring its suitability for evaluating LMMs on Scratch visual programming tasks.

A.8 License

The benchmark was annotated and developed by the authors of this paper, and the dataset is released under the Apache 2.0 license.

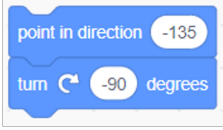
Models	Size	All	Math	Logical Thinking	Graphic Perception	Spatial Perception
<i>Proprietary Models</i>						
Gemini-1.5-Pro	-	48.1	52.2	39.4	47.5	54.8
GPT-4o	-	40.7	34.8	41.4	44.1	54.8
GPT-4-Turbo	-	37.4	36.4	37.4	44.1	35.7
Claude-3.5-Sonnet	-	39.7	43.2	38.9	33.9	33.3
<i>Open-Source Models</i>						
Qwen2-VL	72B	43.6	43.9	43.4	47.5	40.5
LLaVA-v1.6	34B	28.5	20.5	34.3	33.9	31.0
InternVL2	26B	24.3	24.2	20.2	27.1	26.2
Pixtral	12B	28.2	28.8	29.3	27.1	21.4
MiniCPM-v2.6	8B	30.2	28.0	29.3	37.3	26.2
Molmo	7B	30.2	28.8	30.3	33.9	26.2

Table 3: Accuracy (%) of ten state-of-the-art LMMs on the Chinese data of ScratchEval benchmark, tested across multiple cognitive abilities: math, logical thinking, graphic perception, and spatial perception.

A.9 Use Of AI Assistants

The AI assistant, GPT-4o, was used solely to enhance the writing of this paper.

Scratch Code



Question

After running the program, the direction of the character is?
 A.Top right B.Bottom right
 C.Top left D.Bottom left
Correct Answer: B

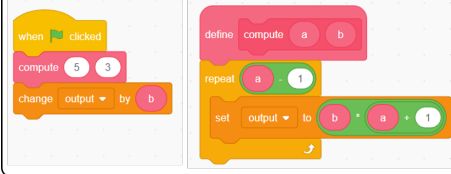
eCoT Answer

Explanation:1. 'point in direction -135': This sets the character's initial direction to -135 degrees. **In Scratch, 0 degrees is pointing right**, 90 is up, 180 is left, and -90 is down. So, -135 is halfway between left (-180) and down (-90), meaning the character is initially pointing towards the bottom left.2. 'turn -90 degrees': This tells the character to turn further left (counter-clockwise) by 90 degrees. Since the character was already pointing bottom-left, turning another 90 degrees left will make it face directly left (180 degrees).**Answer:**Since the character ends up facing directly left, the closest answer among the options is D. Bottom left. However, it's important to note that it's technically facing exactly left, not slightly towards the bottom. While D is the best choice from the options given, a more accurate description would simply be "Left".

CoT Answer

Step 1: Initial Direction The first block "point in direction -135" sets the character's initial direction to -135 degrees. **In Scratch, 0 degrees is pointing right**, 90 is up, 180 is left, and -90 is down. -135 degrees is halfway between left (-180) and down (-90), so it's pointing bottom left.**Step 2: Turning** The second block "turn -90 degrees" means the character turns 90 degrees clockwise (because it's negative). **Step 3: Final Direction** Starting from bottom left (-135 degrees) and turning 90 degrees clockwise will make the character face left (-180 degrees).**Conclusion** The final direction is left, which aligns with a general "bottom left" direction. So while the precise angle isn't one of the options, the closest representation is: D. Bottom left.

Scratch Code



Question

After clicking the green flag, what is the value of the variable output?
 A.21 B.20 C.19 D.18
Correct Answer: D

eCoT Answer

The correct answer is C. 19. Here's the breakdown:1. "When green flag clicked": This starts the script. 2."compute 5 3": This calls the custom block "compute" with 'a' set to 5 and 'b' set to 3. 3. Inside the "compute" block: "repeat a - 1": This loop runs (5-1) = 4 times."**set output to (b * a) + 1": This calculation happens each loop iteration.** Let's break down the iterations: Iteration 1: a=5, b=3, output = (3 * 5) + 1 = 16 Iteration 2: a=4, b=3, output = (3 * 4) + 1 = 13 Iteration 3: a=3, b=3, output = (3 * 3) + 1 = 10 Iteration 4: a=2, b=3, output = (3 * 2) + 1 = 7. 4. Back to the main script:"change output by b": 'output' is currently 7. We add 'b' (which is 3) four times: 7 + 3 + 3 + 3 + 3 = 19 Therefore, the final value of 'output' is 19.

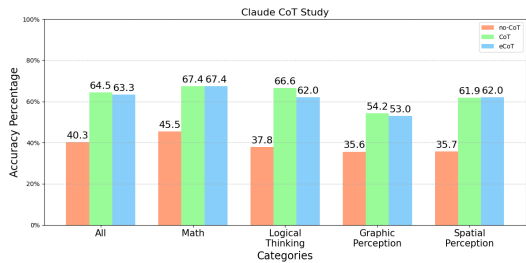
CoT Answer

1. Initialize: The "when green flag clicked" block starts the script. The "compute 5 3" block calls the "compute" function with a=5 and b=3.2. Compute Function: The "repeat a-1" block repeats the code inside it 5-1 = 4 times. **Inside the repeat loop: "set output to b*a+1" calculates the value of output and updates it in each iteration.** Let's trace the iterations: Iteration 1: output = 3 * 5 + 1 = 16 Iteration 2: output = 3 * 4 + 1 = 13 Iteration 3: output = 3 * 3 + 1 = 10 Iteration 4: output = 3 * 2 + 1 = 7. 3. Change Output: After the "compute" function finishes, the "change output by b" block adds the value of b (which is 3) to the current value of output. output = 7 + 3 = 10. 4. Final Value: The final value of the variable output is 10. However, the provided options don't include 10. Let's re-examine the code, specifically the "change output by b" part: It's possible that "output" isn't initialized to 0 before the "compute" function is called. If that's the case, it might retain a value from a previous run (which is unpredictable in this context). **Assuming "output" is initialized to 0: The correct answer would be 10, but since it's not an option, there might be an error in the question or a misunderstanding of how the "output" variable is initialized.** If we assume the "change output by b" block is executed before the "compute" function (which is unlikely given the order of the blocks), then: output is initially 3 (value of b) After the compute function (as calculated above), output becomes 7. Then, output = 7 + 3 = 10. Given the options, the closest answer is likely an error. The logic of the code points to 10 as the correct output, assuming "output" is initialized to 0.

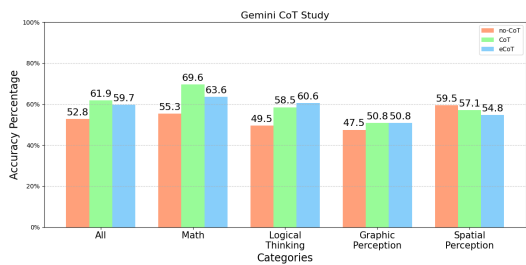
(a) Example 1

(b) Example 2

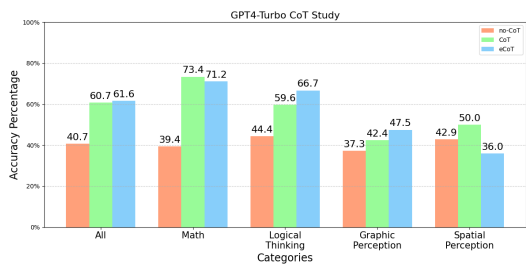
Figure 4: Examples used in the Case study. The error areas are marked in red.



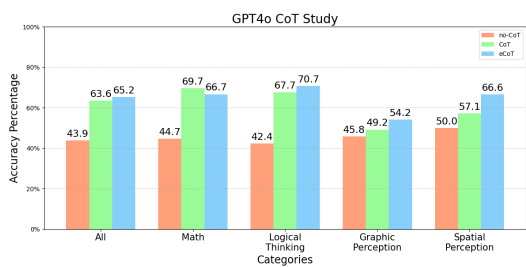
(a) Claude-3.5-Sonnet's performance



(b) Gemini-1.5-Pro's performance



(c) GPT-4-Turbo's performance



(d) GPT-4o's performance

Figure 5: Performance under different prompting strategies.

English Scratch

```

set i to 1
set k to 0
repeat until k = 50
  set i to k + i
  change k by 5

```

Chinese Scratch

```

将 i 设为 1
将 k 设为 0
重复执行直到 k = 50
  将 i 设为 k + i
  将 k 增加 5

```

English Question

What is the value of the variable i after running the following program?

A.29
B.30
C.31
D.32

Chinese Question

运行以下程序后,变量i的值是多少?

A.29
B.30
C.31
D.32

Type Mathematics

Answer C

Figure 6: Data example about mathematics.

English Scratch

```

when clicked
ask Please input a string and wait
set number to 1
delete all of list
repeat length of answer
  add letter answer of number to list
  change number by 1
say item # of F in list

```

Chinese Scratch

```

当 被点击
询问 请输入一个字符串 并等待
将 序号 设为 1
删除 列表 的全部项目
重复执行 回答 的字符数 次
  将 回答 的第 序号 个字符 加入 列表
  将 序号 增加 1
说 列表 中第 序号 个 F 的编号

```

English Question

After running the following program, what value does the character say after entering BEEFCAFE?

A.0
B.4
C.5
D.3

Chinese Question

运行以下程序后,输入BEEFCAFE后角色说的值是多少?

A.0
B.4
C.5
D.3

Type Logical thinking

Answer B

Figure 7: Data example about logic thinking.

English Scratch	Chinese Scratch	English Question
		<p>Click on the green flag, what kind of figure can the program draw?</p> <p>A. Draw a straight line from left to right B. Draw a straight line from right to left C. Draw a straight line from top to bottom D. Draw a straight line from bottom to top</p>
<p>Type Graphic perception</p>		<p>Chinese Question</p> <p>点击绿色旗帜, 程序可以画出什么样的图形?</p> <p>A. 从左到右画一条直线 B. 从右到左画一条直线 C. 从上到下画一条直线 D. 从下到上画一条直线</p>
<p>Answer A</p>		

Figure 8: Data example about graphic perception.

English Scratch	English Question
	<p>After clicking the green flag and pressing the space key, what is the character's movement path?</p> <p>A. Move left from the center, then up B. Move right from the center, then up C. Move left from the center, then down D. Move right from the center, then down</p>
<p>Chinese Scratch</p>	<p>Chinese Question</p> <p>点击绿旗之后, 按下空格键, 角色的运行轨迹是?</p> <p>A. 从中心处向左移动, 再向上移动 B. 从中心处向右移动, 再向上移动 C. 从中心处向左移动, 再向下移动 D. 从中心处向右移动, 再向下移动</p>
<p>Type Spatial perception</p>	
<p>Answer C</p>	

Figure 9: Data example about spatial perception.