



TurtleBench: A Visual Programming Benchmark in Turtle Geometry

Sina Rismanchian, Yasaman Razeghi, Sameer Singh, Shayan Doroudi
University of California, Irvine
{srismanc, yrazeghi, sameer, doroudis}@uci.edu

Abstract

While formal geometric reasoning may be difficult for humans without extensive training, humans seem to have the ability to intuitively reason about geometric patterns in images and scenes from a young age. In contrast, developing large multimodal models (LMMs) capable of similar feats represents a frontier in AI research. We introduce TurtleBench, a benchmark designed to evaluate LMMs' capacity to interpret geometric patterns—given visual examples, textual instructions, or both—and generate precise code outputs. Inspired by turtle geometry, a notion used to teach children foundational coding and geometric concepts, TurtleBench features tasks with patterned shapes that have underlying algorithmic logic. Unlike object detection tasks that typically do not involve understanding underlying patterns, this benchmark combines geometrical reasoning with image understanding. Our evaluation reveals that leading LMMs struggle significantly with these tasks, with GPT-4V achieving only 19% accuracy on the simplest tasks. TurtleBench highlights the gap between human and AI performance in intuitive and visual geometrical understanding, setting the stage for future research in this area.

1 Introduction

Geometric reasoning is a hallmark of human mathematical reasoning that has been studied since the Ancient Greeks. It was a task that attracted early artificial intelligence (AI) researchers and early efforts on building intelligent tutoring systems also focused on geometry. Yet much of the emphasis on geometric reasoning is on axiomatic-deductive geometry. Humans of all ages are naturally good at more intuitive kinds of geometric reasoning that inform how we see and navigate the world. One aspect of this is our ability to look at a geometric shape or complex pattern and construct an algorithm to generate that pattern. We believe this is a powerful task to evaluate large multimodal models (LMMs) for a number of reasons. First of all, constructing patterns in this way reflects an early programming paradigm for teaching kids programming, initially developed in the 1970s with the introduction of the Logo programming language (Papert, 1972, 1980). For several decades, children from a young age have been learning how to procedurally draw geometric patterns and other drawings using code in programming languages like Logo, Scratch, and Python—often as their first introduction to programming. Given LMMs' success in a variety of complex programming tasks, one might expect a programming task that children could solve to be easy. Second, recent research suggests that this ability to procedurally generate shapes may be more fundamental to our psychology than meets the eye. Spelke (2022) claims that from infancy (or even birth), humans have a set of six core knowledge systems, two of which contribute to our understanding of geometry: a *form* system and a *place* system. While the form system allows us to perceive the boundaries of objects, our core knowledge of places interprets geometry in terms of how to navigate an environment (Dillon, 2023). Taking this a step further, Sablé-Meyer et al. (2022) suggest that humans perceive shapes and patterns in terms of procedural programs that could generate them; they demonstrate that the time it takes

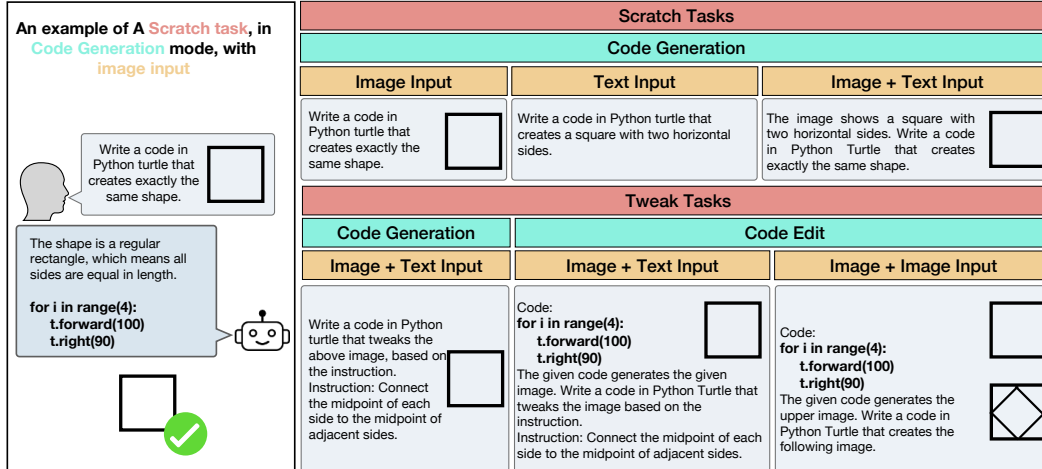


Figure 1: An illustration of existing types and modes in TurtleBench, A task may have a type of Scratch or Tweak, in a mode of code generation or code edit, with various modalities in the input.

for people to process these shapes correlates with the minimum description length of the shape in a Logo-like programming language.

In this work, we introduce TurtleBench, a set of manually crafted image/text to code tasks in turtle geometry (Papert, 1972; Abelson & diSessa, 1986) to evaluate the abilities of these models to combine visual pattern recognition, abstract geometrical reasoning, and Python programming. To ensure the visual inputs and the programming language remain straightforward, TurtleBench harnesses turtle geometry, a concept widely recognized for its effectiveness in introducing programming concepts to children within the K-12 education system. Although turtle programming is now used more as a tool to foster computational thinking, turtle geometry has also been explored as a powerful way of teaching geometry and mathematical reasoning to children (Hoyles & Noss, 1992; Clements & Sarama, 1997). In turtle geometry, a turtle acts as a programmable object that navigates the screen, drawing as it goes and turning at specified angles, to create simple visual patterns. The primary objective within this framework is to generate code capable of producing simple visual inputs. These visual inputs consist of basic geometric shapes, and the programming syntax required is intentionally limited and straightforward. An example of such a task is presented in the left side of Figure 1. As illustrated, the input image is the shape of a simple square and the corresponding code only uses two simple turtle functions (`forward` and `right`) along with a simple `for` loop. This simplicity makes TurtleBench an effective benchmark for evaluating the capabilities of LMMs.

To reflect different real-world use cases of an LMM in the domain of Turtle and also cover the broad range of underlying reasoning abilities, TurtleBench includes 260 tasks with a variety of types and modalities. We conduct an evaluation of leading LMMs on TurtleBench code generation and code editing tasks, utilizing zero-shot and visual chain-of-thought (Singh et al., 2023) approaches across text-only, image-only, and mixed (text and image) input modalities. Our findings reveal that these models generally perform poorly across all setups and variety of tasks and modalities. Our best-performing model, GPT-4V, outperforms Gemini 1.5 Flash yet neither model comes close to solving TurtleBench tasks, as about 75% of the tasks were left completely unsolved. Intriguingly, our results indicate that performance improves when tasks are presented in text, rather than inputting images. This suggests that integrating visual and linguistic information, particularly in domains requiring visual pattern recognition, may need further refinement. All these findings demonstrate that our benchmark poses a challenging task for LMMs, providing valuable insights into their capabilities.

2 Overview of TurtleBench

TurtleBench is a set of 260 tasks that are designed to evaluate LMMs’ performance on vision and language algorithmic reasoning tasks. To ensure the novelty of the tasks and their quality in incorporating authentic geometric shapes and concepts, we craft TurtleBench manually. All the tasks

in TurtleBench are accurately solvable based on the provided information for each, which means that there are no ambiguities or arbitrary parameters leading to inaccuracies in the tasks for humans as well as the models. To remove possible ambiguities in the tasks, two independent annotators worked with us to identify and resolve any unclear instructions. Each task consists of a black-and-white image illustrating a set of abstract geometric shapes as an *input*. An example of this task is presented in Figure 1. TurtleBench is made up of two different types of tasks, these types reflect the methodologies used in turtle geometry to introduce programming to children.

Scratch tasks are intended to show how well a model understands a pattern and translates its understanding to an executable code. In the general case of this type of task, an image is provided, and the requested output is code in Python Turtle that creates the shapes in the image. In all scratch tasks, the model is asked to *generate* the code in Python Turtle for the desired input shape. TurtleBench includes a total of 130 scratch tasks. An example of these tasks is provided in Figure 1, top rows. To distinguish between the models’ visual comprehension and their textual understanding, a subset (31%) of these tasks includes a text description of the image input in addition to the visual representation. This setup facilitates the evaluation of how models respond differently to visual and textual inputs, providing a clearer understanding of their capabilities.

Tweak tasks are intended to measure how well a model uses their understanding of a visual pattern, combined with an instruction to make minimal alterations. Each tweak task presents a model with an image and an instruction; the expected output is Python Turtle code that modifies the shape in the input image according to the given instruction. These tasks are particularly insightful for determining whether a model is merely recalling memorized code for an image, or if it has developed a deeper, more human-like comprehension of the patterns depicted in the images. For instance, a model might be capable of generating code for a certain shape based on training data, but the real challenge lies in its ability to adapt that shape in response to various instructed changes. An example of these tasks is provided in Figure 1, bottom row. Here, the model is given an input image of a rectangle, with an instruction to *connect the midpoint of each side to the midpoint of adjacent sides*. As illustrated in Figure 1, we also introduce a code editing version of the tweak task. In this version, we supply the code corresponding to the input image and then instruct the models to make specific modifications to this code, aiming to achieve a change in the image as per the provided instructions. Detailed information about types of tweaks and their examples is provided in Appendix C.4.

3 Evaluation Setup

In the following section, we evaluate TurtleBench using two state-of-the-art LMMs, GPT-4V and Gemini 1.5 Flash and also an open source model, namely Llava-1.5-13B(Liu et al., 2023) employing greedy decoding in our evaluations. We evaluated two other open models, namely Qwen-VL-Max (Bai et al., 2023) and CogVLM (Wang et al., 2023) on a subset of tasks in TurtleBench. However, CogVLM and Qwen are not successful in producing a syntactically correct Python Turtle piece of code even for the simplest tasks, therefore we limited our benchmark evaluation to the models mentioned above.

We utilize two types of prompting in our experiments, 1) basic, where we simply prompt the model (c.f. Appendix C.2) to do our tasks, and 2) Chain-of-Thought (CoT) prompting (Wei et al., 2022), which has shown to be an effective prompting technique in eliciting reasoning in these models. Specifically, we use a more detailed version of CoT prompting that is tailored to LMMs, namely v-CoT, recently proposed by Singh et al. (2023). The v-CoT approach is inspired by m-CoT (Zhang et al., 2023), which shows higher performance compared to it. This prompting has been shown to improve LMMs’ performance on visual tasks that involved reasoning, such as ARC (Chollet, 2019). This prompt, instructs the model to first extract all the relevant information in the image needed for answering the problem and then to reason step by step based on the information extracted. The specific prompt we used in our experiments is in Appendix C.2

4 Results

4.1 Models perform poorly on TurtleBench

We initially examine the performance of the GPT-4V, Gemini 1.5 Flash and Llava-1.5-13B models on the comprehensive TurtleBench dataset. The findings, detailed in Table 1, reveal a notably poor


	GPT-4V basic	Gemini basic	GPT-4V 0-S CoT	Gemini 0-S CoT	Llava-1.5 basic	Llava-1.5 0-s CoT
<i>Scratch Code Generation</i>						
Image only	16%	7.7%	19.23%	8.46%	1%	1%
<i>Tweak Code Generation</i>						
Image + Text	10%	3.85%	12.3%	7.7%	0%	1%
<i>Tweak Code Edit</i>						
Image + Text	18%	12%	18.46%	18.46%	1%	1%
Image + Image	12%	3%	13.84%	8.46%	NA	NA

Table 1: Performance of GPT-4V, Gemini 1.5 Flash, and Llava-1.5-13B on TurtleBench. Our result shows that models perform poorly on TurtleBench.

performance across the tasks in TurtleBench, with a peak accuracy of 20% achieved by GPT-4V in the *code editing* tasks, facilitated by Chain of Thought (CoT) prompting. In the *scratch* tasks, which represent the simplest problem type within the dataset, GPT-4V’s success rate was just 19%, underscoring the substantial challenges and complexities these tasks pose to the current models. A comparison between CoT and basic prompting within Table 1 illustrates that CoT prompting outperforms basic prompting on the same models, aligning with previous work that indicates CoT enhances models’ reasoning abilities (Zhang et al., 2023). However, despite utilizing CoT prompting, the task remains far from being solved. Additionally, we note a decline in the performance of models when comparing tasks that involve tweaks to those starting from scratch. This observation suggests that models fail to generalize their understanding to tweak tasks, even if they can successfully complete tasks from scratch. Examples of model output in different subsets of the task are provided in Figures 8 and 10.

4.2 Limited Visual Understanding in LLMs: Insights from Textual vs. Visual Tweak Tasks

For tweak tasks, where the AI had to edit existing code, we gave instructions either in natural language or as images (see Figure 1, bottom rows, left two columns). As can be seen by comparing the bottom two rows in Table 1, there is a huge decline in accuracy when instructions were provided visually rather than textually, especially for Gemini. This outcome suggests a disparity in the models’ ability to process visual versus textual instructions, revealing that their reasoning abilities may not align closely with human-like understanding. The assumption that directly viewing the desired outcome simplifies the task contrasts sharply with our findings, highlighting a reliance on textual interpretation for reasoning and a notable limitation in pure visual reasoning capabilities within these models. In Appendix B.3, we provide further evidence of this with additional analyses on scratch tasks by varying the input to those tasks (i.e., visual or textual descriptions).

5 Conclusions

This study introduces TurtleBench, the first of its kind in benchmarks that focus on converting visual inputs to code outputs. The evaluation results from TurtleBench reveal a significant disparity between human capabilities and current state-of-the-art AI models in understanding simple geometric shapes, reasoning about these shapes, and converting such understandings into executable code. This gap underscores the challenges that lie ahead in the quest to enhance AI’s comprehension and problem-solving abilities to match human levels. We believe that TurtleBench serves as a crucial tool in the evaluation of models, offering a clear benchmark that tests the limits of large multimodal models.

References

Harold Abelson and Andrea diSessa. *Turtle geometry: The computer as a medium for exploring mathematics*. MIT press, 1986.

- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pp. 2425–2433, 2015.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 2023.
- Jonas Belouadi, Anne Lauscher, and Steffen Eger. Automatiz: Text-guided synthesis of scientific vector graphics with tikz. *arXiv preprint arXiv:2310.00367*, 2023.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*, 2023.
- Anoop Cherian, Kuan-Chuan Peng, Suhas Lohit, Kevin A Smith, and Joshua B Tenenbaum. Are deep neural networks smarter than second graders? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10834–10844, 2023.
- François Chollet. On the measure of intelligence. *arXiv preprint arXiv:1911.01547*, 2019.
- Douglas H Clements and Julie Sarama. Children’s mathematical reasoning with the turtle programming metaphor. In *Mathematical Reasoning*, pp. 313–337. Routledge, 1997.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning, 2023.
- Moira R Dillon. Divisive language, 2023.
- Kevin Ellis, Lionel Wong, Maxwell Nye, Mathias Sable-Meyer, Luc Cary, Lore Anaya Pozo, Luke Hewitt, Armando Solar-Lezama, and Joshua B Tenenbaum. Dreamcoder: growing generalizable, interpretable knowledge with wake–sleep bayesian program learning. *Philosophical Transactions of the Royal Society A*, 381(2251):20220050, 2023.
- Kenneth Forbus, Jeffrey Usher, Andrew Lovett, Kate Lockwood, and Jon Wetzel. Cogsketch: Sketch understanding for cognitive science research and for education. *Topics in Cognitive Science*, 3(4): 648–666, 2011.
- Gabriel Grand, Lionel Wong, Matthew Bowers, Theo X Olausson, Muxin Liu, Joshua B Tenenbaum, and Jacob Andreas. Lilo: Learning interpretable libraries by compressing and documenting code. *arXiv preprint arXiv:2310.19791*, 2023.
- Celia Hoyles and Richard Noss. *Learning mathematics and Logo*. MIT Press, 1992.
- Huaizu Jiang, Xiaojian Ma, Weili Nie, Zhiding Yu, Yuke Zhu, and Anima Anandkumar. Bongard-hoi: Benchmarking few-shot visual reasoning for human-object interactions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 19056–19065, 2022.
- Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2901–2910, 2017.
- Brenden M Lake and Steven T Piantadosi. People infer recursive visual concepts from just a few examples. *Computational Brain & Behavior*, 3(1):54–65, 2020.
- Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- Yi Lin and Moira R Dillon. We are wanderers: Abstract geometry reflects spatial navigation. *Journal of Experimental Psychology: General*, 2023.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning, 2023.

Raja Marjeh, Pol Van Rijn, Ilya Sucholutsky, Theodore R Sumers, Harin Lee, Thomas L Griffiths, and Nori Jacoby. Words are all you need? language as an approximation for human similarity judgments. *arXiv preprint arXiv:2206.04105*, 2022.

Melanie Mitchell, Alessandro B Palmarini, and Arseny Moskvichev. Comparing humans, gpt-4, and gpt-4v on abstraction and reasoning tasks. *arXiv preprint arXiv:2311.09247*, 2023.

Weili Nie, Zhiding Yu, Lei Mao, Ankit B Patel, Yuke Zhu, and Anima Anandkumar. Bongard-logo: A new benchmark for human-level concept learning and reasoning. *Advances in Neural Information Processing Systems*, 33:16468–16480, 2020.

OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. Gpt-4 technical report, 2024.

Seymour Papert. On making a theorem for a child. In *Proceedings of the ACM annual conference-Volume 1*, pp. 345–349, 1972.

- Seymour Papert. *Mindstorms: Children, computers, and powerful ideas*. Basic Books, Inc., 1980.
- Joshua S Rule, Joshua B Tenenbaum, and Steven T Piantadosi. The child as hacker. *Trends in cognitive sciences*, 24(11):900–915, 2020.
- Mathias Sablé-Meyer, Kevin Ellis, Josh Tenenbaum, and Stanislas Dehaene. A language of thought for the mental representation of geometric shapes. *Cognitive Psychology*, 139:101527, 2022. ISSN 0010-0285. doi: <https://doi.org/10.1016/j.cogpsych.2022.101527>. URL <https://www.sciencedirect.com/science/article/pii/S0010028522000639>.
- Mukul Singh, José Cambronero, Sumit Gulwani, Vu Le, and Gust Verbruggen. Assessing gpt4-v on structured reasoning tasks. *arXiv preprint arXiv:2312.11524*, 2023.
- Elizabeth S Spelke. *What babies know: Core knowledge and composition volume 1*, volume 1. Oxford University Press, 2022.
- Elizabeth S Spelke and Katherine D Kinzler. Core knowledge. *Developmental science*, 10(1):89–96, 2007.
- Iliia Sucholutsky and Tom Griffiths. Alignment with human representations supports robust few-shot learning. *Advances in Neural Information Processing Systems*, 36, 2024.
- Gemini Team, 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 Lillicrap, Angeliki Lazaridou, Orhan Firat, James Molloy, Michael Isard, Paul R. 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 Piqueras, 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, Jakub Sygnowski, Alexandre Frechette, Charlotte Smith, Laura Culp, Lev Proleev, Yi Luan, Xi Chen, James Lottes, Nathan Schucher, Federico Lebron, Alban Rrustemi, Natalie Clay, Phil Crone, Tomas Kocisky, Jeffrey Zhao, Bartek Perz, Dian Yu, Heidi Howard, Adam Bloniarz, Jack W. Rae, Han Lu, Laurent Sifre, Marcello Maggioni, Fred Alcober, Dan Garrette, Megan Barnes, Shantanu Thakoor, Jacob Austin, Gabriel Barth-Maron, William Wong, Rishabh Joshi, Rahma Chaabouni, Deeni Fatiha, Arun Ahuja, Ruibo Liu, Yunxuan Li, Sarah Cogan, Jeremy Chen, Chao Jia, Chenjie Gu, Qiao Zhang, Jordan Grimstad, Ale Jakse Hartman, Martin Chadwick, Gaurav Singh Tomar, Xavier Garcia, Evan Senter, Emanuel Taropa, Thanumalayan Sankaranarayanan Pillai, Jacob Devlin, Michael Laskin, Diego de Las Casas, Dasha Valter, Connie Tao, Lorenzo Blanco, Adrià Puigdomènech Badia, David Reitter, Mianna Chen, Jenny Brennan, Clara Rivera, Sergey Brin, Shariq Iqbal, Gabriela Surita, Jane Labanowski, Abhi Rao, Stephanie Winkler, Emilio Parisotto, Yiming Gu, Kate Olszewska, Yujing Zhang, Ravi Addanki, Antoine Miech, Annie Louis, Laurent El Shafey, Denis Teplyashin, Geoff Brown, Elliot Catt, Nithya Attaluri, Jan Balaguer, Jackie Xiang, Pidong Wang, Zoe Ashwood, Anton Briukhov, Albert Webson, Sanjay Ganapathy, Smit Sanghavi, Ajay Kannan, Ming-Wei Chang, Axel Stjerngren, Josip Djolonga, Yuting Sun, Ankur Bapna, Matthew Aitchison, Pedram Pejman, Henryk Michalewski, Tianhe Yu, Cindy Wang, Juliette Love, Junwhan Ahn, Dawn Bloxwich, Kehang Han, Peter Humphreys, Thibault Sellam, James Bradbury, Varun Godbole, Sina Samangooei, Bogdan Damoc, Alex Kaskasoli, Sébastien M. R. Arnold, Vijay Vasudevan, Shubham Agrawal, Jason Riesa, Dmitry Lepikhin, Richard Tanburn, Srivatsan Srinivasan, Hyeontaek Lim, Sarah Hodgkinson, Pranav Shyam, Johan Ferret, Steven Hand, Ankush Garg, Tom Le Paine, Jian Li, Yujia Li, Minh Giang, Alexander Neitz, Zaheer Abbas, Sarah York, Machel Reid, Elizabeth Cole, Aakanksha Chowdhery, Dipanjan Das, Dominika Rogozińska, Vitaly Nikolaev, Pablo Sprechmann, Zachary Nado, Lukas Zilka, Flavien Prost, Luheng He, Marianne Monteiro, Gaurav Mishra, Chris Welty, Josh Newlan, Dawei Jia, Miltiadis Allamanis, Clara Huiyi Hu, Raoul de Liedekerke, Justin Gilmer, Carl Saroufim, Shruti Rijhwani, Shaobo Hou, Disha Shrivastava, Anirudh Baddepudi, Alex Goldin, Adnan Ozturk, Albin Cassirer, Yunhan Xu, Daniel Sohn, Devendra Sachan, Reinald Kim Amplayo, Craig Swanson, Dessie Petrova, Shashi Narayan, Arthur Guez, Siddhartha Brahma, Jessica Landon, Miteyan Patel, Ruizhe Zhao, Kevin Villela, Luyu Wang, Wenhao Jia, Matthew Rahtz, Mai Giménez, Legg Yeung, Hanzhao Lin, James Keeling, Petko Georgiev, Diana Mincu, Boxi Wu, Salem Haykal, Rachel Saputro, Kiran Vodrahalli, James

Qin, Zeynep Cankara, Abhanshu Sharma, Nick Fernando, Will Hawkins, Behnam Neyshabur, Solomon Kim, Adrian Hutter, Priyanka Agrawal, Alex Castro-Ros, George van den Driessche, Tao Wang, Fan Yang, Shuo yin Chang, Paul Komarek, Ross McIlroy, Mario Lučić, Guodong Zhang, Wael Farhan, Michael Sharman, Paul Natsev, Paul Michel, Yong Cheng, Yamini Bansal, Siyuan Qiao, Kris Cao, Siamak Shakeri, Christina Butterfield, Justin Chung, Paul Kishan Rubenstein, Shivani Agrawal, Arthur Mensch, Kedar Soparkar, Karel Lenc, Timothy Chung, Aedan Pope, Loren Maggiore, Jackie Kay, Priya Jhakra, Shibo Wang, Joshua Maynez, Mary Phuong, Taylor Tobin, Andrea Tacchetti, Maja Trebacz, Kevin Robinson, Yash Katariya, Sebastian Riedel, Paige Bailey, Kefan Xiao, Nimesh Ghelani, Lora Aroyo, Ambrose Slone, Neil Houlsby, Xuehan Xiong, Zhen Yang, Elena Gribovskaya, Jonas Adler, Mateo Wirth, Lisa Lee, Music Li, Thais Kagohara, Jay Pavagadhi, Sophie Bridgers, Anna Bortsova, Sanjay Ghemawat, Zafarali Ahmed, Tianqi Liu, Richard Powell, Vijay Bolina, Mariko Inuma, Polina Zablotskaia, James Besley, Da-Woon Chung, Timothy Dozat, Ramona Comanescu, Xiance Si, Jeremy Greer, Guolong Su, Martin Polacek, Raphaël Lopez Kaufman, Simon Tokumine, Hexiang Hu, Elena Buchatskaya, Yingjie Miao, Mohamed Elhawaty, Aditya Siddhant, Nenad Tomasev, Jinwei Xing, Christina Greer, Helen Miller, Shereen Ashraf, Aurko Roy, Zizhao Zhang, Ada Ma, Angelos Filos, Milos Besta, Rory Blevins, Ted Klimenko, Chih-Kuan Yeh, Soravit Changpinyo, Jiaqi Mu, Oscar Chang, Mantas Pajarskas, Carrie Muir, Vered Cohen, Charline Le Lan, Krishna Haridasan, Amit Marathe, Steven Hansen, Sholto Douglas, Rajkumar Samuel, Mingqiu Wang, Sophia Austin, Chang Lan, Jiepu Jiang, Justin Chiu, Jaime Alonso Lorenzo, Lars Lowe Sjösund, Sébastien Cevey, Zach Gleicher, Thi Avrahami, Anudhyan Boral, Hansa Srinivasan, Vittorio Selo, Rhys May, Konstantinos Aisopos, Léonard Hussenot, Livio Baldini Soares, Kate Baumli, Michael B. Chang, Adrià Recasens, Ben Caine, Alexander Pritzel, Filip Pavetic, Fabio Pardo, Anita Gergely, Justin Frye, Vinay Ramasesh, Dan Horgan, Kartikeya Badola, Nora Kassner, Subhrajit Roy, Ethan Dyer, Víctor Campos, Alex Tomala, Yunhao Tang, Dalia El Badawy, Elspeth White, Basil Mustafa, Oran Lang, Abhishek Jindal, Sharad Vikram, Zhitao Gong, Sergi Caelles, Ross Hemsley, Gregory Thornton, Fangxiaoyu Feng, Wojciech Stokowiec, Ce Zheng, Phoebe Thacker, Çağlar Ünlü, Zhishuai Zhang, Mohammad Saleh, James Svensson, Max Bileschi, Piyush Patil, Ankesh Anand, Roman Ring, Katerina Tsihlias, Arpi Vezer, Marco Selvi, Toby Shevlane, Mikel Rodriguez, Tom Kwiatkowski, Samira Daruki, Keran Rong, Allan Dafoe, Nicholas FitzGerald, Keren Gu-Lemberg, Mina Khan, Lisa Anne Hendricks, Marie Pellat, Vladimir Feinberg, James Cobon-Kerr, Tara Sainath, Maribeth Rauh, Sayed Hadi Hashemi, Richard Ives, Yana Hasson, YaGuang Li, Eric Noland, Yuan Cao, Nathan Byrd, Le Hou, Qingze Wang, Thibault Sottiaux, Michela Paganini, Jean-Baptiste Lespiau, Alexandre Moufarek, Samer Hassan, Kaushik Shivakumar, Joost van Amersfoort, Amol Mandhane, Pratik Joshi, Anirudh Goyal, Matthew Tung, Andrew Brock, Hannah Sheahan, Vedant Misra, Cheng Li, Nemanja Rakićević, Mostafa Dehghani, Fangyu Liu, Sid Mittal, Junhyuk Oh, Seb Noury, Eren Sezener, Fantine Huot, Matthew Lamm, Nicola De Cao, Charlie Chen, Gamaleldin Elsayed, Ed Chi, Mahdis Mahdieh, Ian Tenney, Nan Hua, Ivan Petrychenko, Patrick Kane, Dylan Scandinaro, Rishub Jain, Jonathan Uesato, Romina Datta, Adam Sadovsky, Oskar Bunyan, Dominik Rabiej, Shimu Wu, John Zhang, Gautam Vasudevan, Edouard Leurent, Mahmoud Alnahlawi, Ionut Georgescu, Nan Wei, Ivy Zheng, Betty Chan, Pam G Rabinovitch, Piotr Stanczyk, Ye Zhang, David Steiner, Subhajit Naskar, Michael Azzam, Matthew Johnson, Adam Paszke, Chung-Cheng Chiu, Jaume Sanchez Elias, Afroz Mohiuddin, Faizan Muhammad, Jin Miao, Andrew Lee, Nino Vieillard, Sahitya Potluri, Jane Park, Elnaz Davoodi, Jiageng Zhang, Jeff Stanway, Drew Garmon, Abhijit Karmarkar, Zhe Dong, Jong Lee, Aviral Kumar, Luowei Zhou, Jonathan Evens, William Isaac, Zhe Chen, Johnson Jia, Anselm Levskaya, Zhenkai Zhu, Chris Gorgolewski, Peter Grabowski, Yu Mao, Alberto Magni, Kaisheng Yao, Javier Snaider, Norman Casagrande, Paul Suganthan, Evan Palmer, Geoffrey Irving, Edward Loper, Manaal Faruqui, Isha Arkatkar, Nanxin Chen, Izhak Shafran, Michael Fink, Alfonso Castaño, Irene Giannoumis, Wooyeol Kim, Mikolaj Rybiński, Ashwin Sreevatsa, Jennifer Prendki, David Soergel, Adrian Goedeckemeyer, Willi Gierke, Mohsen Jafari, Meenu Gaba, Jeremy Wiesner, Diana Gage Wright, Yawen Wei, Harsha Vashisht, Yana Kulizhskaya, Jay Hoover, Maigo Le, Lu Li, Chimezie Iwuanyanwu, Lu Liu, Kevin Ramirez, Andrey Khorlin, Albert Cui, Tian LIN, Marin Georgiev, Marcus Wu, Ricardo Aguilar, Keith Pallo, Abhishek Chakladar, Alena Repina, Xihui Wu, Tom van der Weide, Priya Ponnappalli, Caroline Kaplan, Jiri Simsa, Shuangfeng Li, Olivier Dousse, Fan Yang, Jeff Piper, Nathan Ie, Minnie Lui, Rama Pasumarthi, Nathan Lintz, Anitha Vijayakumar, Lam Nguyen Thiet, Daniel Andor, Pedro Valenzuela, Cosmin Paduraru, Daiyi Peng, Katherine Lee, Shuyuan Zhang, Somer Greene, Duc Dung Nguyen, Paula Kurylowicz, Sarmishta Velury, Sebastian Krause, Cassidy Hardin, Lucas Dixon, Lili Janzer, Kiam Choo, Ziqiang Feng, Biao Zhang, Achintya Singhal, Tejasi Latkar, Mingyang Zhang, Quoc Le,

Elena Allica Abellan, Dayou Du, Dan McKinnon, Natasha Antropova, Tolga Bolukbasi, Orgad Keller, David Reid, Daniel Finchelstein, Maria Abi Raad, Remi Crocker, Peter Hawkins, Robert Dadashi, Colin Gaffney, Sid Lall, Ken Franko, Egor Filonov, Anna Bulanova, Rémi Leblond, Vikas Yadav, Shirley Chung, Harry Askham, Luis C. Cobo, Kelvin Xu, Felix Fischer, Jun Xu, Christina Sorokin, Chris Alberti, Chu-Cheng Lin, Colin Evans, Hao Zhou, Alek Dimitriev, Hannah Forbes, Dylan Banarse, Zora Tung, Jeremiah Liu, Mark Omernick, Colton Bishop, Chintu Kumar, Rachel Sterneck, Ryan Foley, Rohan Jain, Swaroop Mishra, Jiawei Xia, Taylor Bos, Geoffrey Cideron, Ehsan Amid, Francesco Piccinno, Xingyu Wang, Praseem Banzal, Petru Gurita, Hila Noga, Premal Shah, Daniel J. Mankowitz, Alex Polozov, Nate Kushman, Victoria Krakovna, Sasha Brown, MohammadHossein Bateni, Dennis Duan, Vlad Firoiu, Meghana Thotakuri, Tom Natan, Anhad Mohananey, Matthieu Geist, Sidharth Mudgal, Sertan Girgin, Hui Li, Jiayu Ye, Ofir Roval, Reiko Tojo, Michael Kwong, James Lee-Thorp, Christopher Yew, Quan Yuan, Sumit Bagri, Danila Sinopalnikov, Sabela Ramos, John Mellor, Abhishek Sharma, Aliaksei Severyn, Jonathan Lai, Kathy Wu, Heng-Tze Cheng, David Miller, Nicolas Sonnerat, Denis Vnukov, Rory Greig, Jennifer Beattie, Emily Caveness, Libin Bai, Julian Eisenschlos, Alex Korchemniy, Tomy Tsai, Mimi Jasarevic, Weize Kong, Phuong Dao, Zeyu Zheng, Frederick Liu, Fan Yang, Rui Zhu, Mark Geller, Tian Huey Teh, Jason Sanmiya, Evgeny Gladchenko, Nejc Trdin, Andrei Sozanschi, Daniel Toyama, Evan Rosen, Sasan Tavakkol, Linting Xue, Chen Elkind, Oliver Woodman, John Carpenter, George Papamakarios, Rupert Kemp, Sushant Kafle, Tanya Grunina, Rishika Sinha, Alice Talbert, Abhimanyu Goyal, Diane Wu, Denese Owusu-Afryie, Cosmo Du, Chloe Thornton, Jordi Pont-Tuset, Pradyumna Narayana, Jing Li, Sabaer Fatehi, John Wieting, Omar Ajmeri, Benigno Uria, Tao Zhu, Yeongil Ko, Laura Knight, Amélie Héliou, Ning Niu, Shane Gu, Chenxi Pang, Dustin Tran, Yeqing Li, Nir Levine, Ariel Stolovich, Norbert Kalb, Rebeca Santamaria-Fernandez, Sonam Goenka, Wenny Yustalim, Robin Strudel, Ali Elqursh, Balaji Lakshminarayanan, Charlie Deck, Shyam Upadhyay, Hyo Lee, Mike Dusenberry, Zonglin Li, Xuezhi Wang, Kyle Levin, Raphael Hoffmann, Dan Holtmann-Rice, Olivier Bachem, Summer Yue, Sho Arora, Eric Malmi, Daniil Mirylenka, Qijun Tan, Christy Koh, Soheil Hassas Yeganeh, Siim Pöder, Steven Zheng, Francesco Pongetti, Mukarram Tariq, Yanhua Sun, Lucian Ionita, Mojtaba Seyedhosseini, Pouya Tafti, Ragha Kotikalapudi, Zhiyu Liu, Anmol Gulati, Jasmine Liu, Xinyu Ye, Bart Chrzaszcz, Lily Wang, Nikhil Sethi, Tianrun Li, Ben Brown, Shreya Singh, Wei Fan, Aaron Parisi, Joe Stanton, Chenkai Kuang, Vinod Koverkathu, Christopher A. Choquette-Choo, Yunjie Li, TJ Lu, Abe Ittycheriah, Prakash Shroff, Pei Sun, Mani Varadarajan, Sanaz Bahargam, Rob Willoughby, David Gaddy, Ishita Dasgupta, Guillaume Desjardins, Marco Cornero, Brona Robenek, Bhavishya Mittal, Ben Albrecht, Ashish Shenoy, Fedor Moiseev, Henrik Jacobsson, Alireza Ghaffarkhah, Morgane Rivière, Alanna Walton, Clément Crepy, Alicia Parrish, Yuan Liu, Zongwei Zhou, Clement Farabet, Carey Radebaugh, Praveen Srinivasan, Claudia van der Salm, Andreas Fijeldand, Salvatore Scellato, Eri Latorre-Chimoto, Hanna Klimczak-Plucińska, David Bridson, Dario de Cesare, Tom Hudson, Piermaria Mendolicchio, Lexi Walker, Alex Morris, Ivo Penchev, Matthew Mauger, Alexey Guseynov, Alison Reid, Seth Odoom, Lucia Loher, Victor Cotruta, Madhavi Yenugula, Dominik Grewe, Anastasia Petrushkina, Tom Duerig, Antonio Sanchez, Steve Yadlowsky, Amy Shen, Amir Globerson, Adam Kurzrok, Lynette Webb, Sahil Dua, Dong Li, Preethi Lahoti, Surya Bhupatiraju, Dan Hurt, Haroon Qureshi, Ananth Agarwal, Tomer Shani, Matan Eyal, Anuj Khare, Shreyas Rammohan Belle, Lei Wang, Chetan Tekur, Mihir Sanjay Kale, Jinliang Wei, Ruoxin Sang, Brennan Saeta, Tyler Liechty, Yi Sun, Yao Zhao, Stephan Lee, Pandu Nayak, Doug Fritz, Manish Reddy Vuyyuru, John Aslanides, Nidhi Vyas, Martin Wicke, Xiao Ma, Taylan Bilal, Evgenii Eltyshev, Daniel Balle, Nina Martin, Hardie Cate, James Manyika, Keyvan Amiri, Yelin Kim, Xi Xiong, Kai Kang, Florian Luisier, Nilesh Tripuraneni, David Madras, Mandy Guo, Austin Waters, Oliver Wang, Joshua Ainslie, Jason Baldridge, Han Zhang, Garima Pruthi, Jakob Bauer, Feng Yang, Riham Mansour, Jason Gelman, Yang Xu, George Polovets, Ji Liu, Honglong Cai, Warren Chen, XiangHai Sheng, Emily Xue, Sherjil Ozair, Adams Yu, Christof Angermueller, Xiaowei Li, Weiren Wang, Julia Wiesinger, Emmanouil Koukoumidis, Yuan Tian, Anand Iyer, Madhu Gurusurthy, Mark Goldenson, Parashar Shah, MK Blake, Hongkun Yu, Anthony Urbanowicz, Jennimaria Palomaki, Chrisantha Fernando, Kevin Brooks, Ken Durden, Harsh Mehta, Nikola Momchev, Elahe Rahimtoroghi, Maria Georgaki, Amit Raul, Sebastian Ruder, Morgan Redshaw, Jinhyuk Lee, Komal Jalan, Dinghua Li, Ginger Perng, Blake Hechtman, Parker Schuh, Milad Nasr, Mia Chen, Kieran Milan, Vladimir Mikulik, Trevor Strohman, Juliana Franco, Tim Green, Demis Hassabis, Koray Kavukcuoglu, Jeffrey Dean, and Oriol Vinyals. Gemini: A family of highly capable multimodal models, 2023.

- Weihan Wang, Qingsong Lv, Wenmeng Yu, Wenyi Hong, Ji Qi, Yan Wang, Junhui Ji, Zhuoyi Yang, Lei Zhao, Xixuan Song, et al. Cogvlm: Visual expert for pretrained language models. *arXiv preprint arXiv:2311.03079*, 2023.
- Weihan Wang, Qingsong Lv, Wenmeng Yu, Wenyi Hong, Ji Qi, Yan Wang, Junhui Ji, Zhuoyi Yang, Lei Zhao, Xixuan Song, Jiazheng Xu, Bin Xu, Juanzi Li, Yuxiao Dong, Ming Ding, and Jie Tang. Cogvlm: Visual expert for pretrained language models, 2024.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.
- Catherine Wong, Kevin M Ellis, Joshua Tenenbaum, and Jacob Andreas. Leveraging language to learn program abstractions and search heuristics. In *International conference on machine learning*, pp. 11193–11204. PMLR, 2021.
- Xiangyu Wu, Yang Yang, Shengdong Xu, Yifeng Wu, Qingguo Chen, and Jianfeng Lu. Solution for smart-101 challenge of iccv multi-modal algorithmic reasoning task 2023. *arXiv preprint arXiv:2310.06440*, 2023.
- Chi Zhang, Feng Gao, Baoxiong Jia, Yixin Zhu, and Song-Chun Zhu. Raven: A dataset for relational and analogical visual reasoning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 5317–5327, 2019.
- Zhuosheng Zhang, Aston Zhang, Mu Li, Hai Zhao, George Karypis, and Alex Smola. Multimodal chain-of-thought reasoning in language models. *arXiv preprint arXiv:2302.00923*, 2023.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt-4: Enhancing vision-language understanding with advanced large language models, 2023.

A Related Work

A.1 Large Multi-modal Models

Recent advancements in foundational multimodal models have marked a significant stride towards developing generalist AI systems capable of understanding and integrating information across different modalities to solve tasks without the need for task-specific fine-tuning. Among these models are closed source models such as Gemini 1.5 Flash (Team et al., 2023), GPT-4V (OpenAI et al., 2024), and open source models as LLaVA-1.5 (Liu et al., 2023), Mini-GPT4 (Zhu et al., 2023), InstructBLIP (Dai et al., 2023) and CogVLM (Wang et al., 2024). The versatility and multimodal understanding exhibited by these foundational multimodal models have positioned them as prime candidates for applications such as AI software engineers or programming tutors for children. Our work evaluates the efficacy of these popular models on image/text-to-code tasks, measuring their potential in vision/programming context.

A.2 Probabilistic Program Induction

Recent work in Bayesian cognitive science has modeled various aspects of cognition and learning as probabilistic program induction (Lake et al., 2015; Lake & Piantadosi, 2020; Rule et al., 2020; Ellis et al., 2023; Wong et al., 2021; Grand et al., 2023). This has involved both modeling human cognition as program induction as well as designing machine learning algorithms that can generate programs for various tasks, including the kind of turtle geometry task we study here. Ellis et al. (2023) developed the DreamCoder algorithm which can learn to induce programs by using self-supervision to incrementally build up a library of programs and train a neural network to search to find the best program for a given task. They created a dataset of 160 turtle programming tasks. In contrast to our approach, where we assess the performance of out-of-the-box LMMs, DreamCoder is trained on a training set of images (i.e., half of the dataset). However, it is interesting that the algorithm is trained in an unsupervised fashion; that is, DreamCoder never receives the code used to generate the images and learns that from experience. Wong et al. (2021) extended this work by developing an algorithm (LAPS) that can induce programs given both the task and linguistic annotations for the task. They

used a dataset of 311 turtle graphics with greater complexity than the original DreamCoder dataset. While their dataset includes linguistic annotations, their dataset does not include tweak tasks like in TurtleBench. Additionally, their tasks often include arbitrary aspects (for example, a gap with unspecified distance between two shapes) that makes evaluation hard; in our tasks, the positional relationships between shapes should be easy to infer exactly and hence we can evaluate models by comparing exactly with ground truth shapes. Moreover, neither of these datasets have been framed as a benchmark for visual program induction and have not been considered for evaluating LMMs. Perhaps the approach closest to our work is by Grand et al. (2023), who combined LLMs with a symbolic program induction algorithm and evaluated the performance of their model (LILO) on the turtle geometry task using the aforementioned dataset. Averaged over several runs, the performance of the best versions of these approaches on the turtle geometry task is as follows: 43% for DreamCoder, 82% for LAPS, 49% for LILO, and 32% for a LLM solver. These results seem to suggest that probabilistic programming approaches (such as LAPS) can greatly outperform LMMs on visual programming tasks. We note that the performance of the LLM solver (32%) is comparable to the performance of GPT-4V on our text-only input (37%; see Table 4). Future work could assess the performance of probabilistic program induction methods like LAPS on TurtleBench.

A.3 Mutimodal Algorithmic Reasoning

The existing literature features a range of studies that evaluate these models using naturalistic images (Jiang et al., 2022; Johnson et al., 2017; Antol et al., 2015), yet humans naturally are able to reason over abstract shapes (Chollet, 2019; Zhang et al., 2019; Spelke & Kinzler, 2007) and also many use cases of LMMs involve understanding abstract shapes and sketches (Forbus et al., 2011; Nie et al., 2020). Moreover, unlike naturalistic images (Marjeh et al., 2022; Sucholutsky & Griffiths, 2024), the relationship between language and abstract shapes is highly intertwined as minimal alterations in language can lead to different visual perceptions in humans (Dillon, 2023; Lin & Dillon, 2023). The Multimodal Algorithmic Reasoning (MAR) task tests multi-modal models on fundamental skills understandable by children, focusing on interpreting visual and linguistic information to answer questions. Perhaps the most relevant work to ours is the paper by Cherian et al. (2023) in which they introduced a dataset with 101 multiple-choice questions inspired by the Math Kangaroo contest for 6 to 8-year-olds, involving images and texts that the model must analyze together. The task has been shown to be challenging for multimodal deep neural networks, and the following trials to solve the problem have gained less than 25% accuracy on the private test set Wu et al. (2023). Our proposed benchmark pushes the evaluation of LMMs forward as TurtleBench includes abstract geometric shapes, and the task only relies on knowledge and reasoning over a set of simple functions in the Python Turtle library. The open-ended nature of our benchmark and its flexibility over different modalities makes evaluating different aspects of vision and language algorithmic reasoning in the models more reliable.

B Additional Analyses

B.1 Models fail to generalize

Given that these models have been extensively trained on vast datasets sourced from the internet, there’s an underlying uncertainty regarding the source of their performance—albeit poor—on the TurtleBench tasks. Specifically, it remains unclear whether this performance is the result of the models’ ability to memorize aspects of our tasks, rather than genuinely understanding and solving them based on their programming and reasoning capabilities. To address this issue, our next step is to evaluate the true generalization ability of these models. By doing so, we aim to distinguish between superficial learning, potentially influenced by memorization, and genuine comprehension and problem-solving skills. To measure the generalizability of the model’s performance, we define an arbitrary set of commands based on the turtle module in Python. In other words, we developed a class called Rabbit that inherits the Turtle class from the turtle module. Although the functions of the Rabbit class are functionally identical to those in the original turtle module, they are nominally distinct. This differentiation allows us to evaluate the models’ ability to apply their knowledge to unfamiliar yet equivalent command sets. The definition of the Rabbit class in Python is provided in Appendix C.3.2. We perform a zero-shot CoT prompting to elicit the code using the new set of





	 GPT-4V Turtle CoT	 GPT-4V Rabbit CoT	 Gemini Turtle CoT	 Gemini Rabbit CoT
<i>Scratch Code Generation</i>				
Image only Input	19%	6%	8.46%	3%
<i>Tweak Code Generation</i>				
Image + Text	12%	2%	7.7%	1%

Table 2: Performance of GPT-4V and Gemini 1.5 Flash on generalization tasks, in these tasks, we defined Rabbit, a new set of functions practically equivalent to but nominally different from the ones in Python Turtle. The performance in Rabbit drastically drops, showing poor generalization abilities in both models.


	 Python Turtle Output	Any Output
<i>Scratch Code Generation</i>		
Image only Input	19.23%	21.6%
<i>Tweak Code Generation</i>		
Image + Text	12.3%	15.1%

Table 3: Performance of (CoT) prompting with GPT-4V on tasks involving code generation for simple geometric shapes in any programming language of the model’s choice reveals that models struggle significantly, even in their preferred programming language.

commands. In the context window, we provide a verbal definition of each function in the Rabbit class. The results of comparing the models’ performance using the Rabbit class versus the standard Python Turtle module are presented in Table 2. We observe that, although both models were capable of generating executable pieces of code with the new class, there is a huge decline in their performance relative to their performance with the conventional Python Turtle module. This finding suggests that the visual reasoning in these models is not robust to syntax changes, and it is likely that they rely on training memorization rather than pure reasoning.

B.2 Assessing Model Proficiency Across Programming Languages

The initial suspicion might be that the models struggle with tasks in turtle geometry due to a lack of exposure to specific programming syntax during pretraining. However, to investigate whether the challenge lies not in syntax familiarity but in understanding visual input and translating this understanding into effective programming, we modify our approach with GPT-4V. We choose GPT-4V as it is our best-performing model in the main task. We allow it to generate code using any library, language, or similar tools it deems appropriate, such as Matplotlib, TikZ, etc., without restricting it to the Python Turtle library. The prompt for this subset of tasks is presented in Appendix C.2.4. We manually evaluate the GPT-4V output for this task. Despite this freedom, we observe no significant improvement in performance. The model chooses Matplotlib for 50% of the tasks and offers pseudocode for 2%, with the remainder reverting to Python Turtle, even though we do not specify Python Turtle in the prompts. Notably, it avoids using TikZ, despite its mention in the prompt and proven capabilities in prior work to produce TikZ code (Bubeck et al., 2023; Belouadi et al., 2023). This outcome underscores a deeper issue than syntax familiarity: the models’ fundamental challenge is accurately interpreting visual input and applying this understanding to generate corresponding programming code.

B.3 Limited Visual Understanding in LMMs: Insights from Scratch Tasks

One of the questions regarding LMMs’ abilities in visual abstraction and understanding tasks is the extent the incorporation of the visual component has enhanced their abilities in reasoning (Mitchell et al., 2023). In resonance with what Mitchell et al. (2023) found, here we also found that the vision component contributes poorly to fostering the models’ visual reasoning abilities, at least in the domain of TurtleBench. We explored this in the context of tweak tasks in Section 4.2. Here, we explore it in

	GPT-4V basic	Gemini basic	GPT-4V CoT	Gemini CoT
<i>Scratch Code Generation</i>				
Image only Input	26%	7.7%	29%	8.46%
Text only Input	37%	25.1%	38%	18.51%
Image and Text Input	38%	22.2%	40%	22.22%

Table 4: Performance of GPT-4V and Gemini 1.5 Flash on TurtleBench, for comparing visual vs. text input on Scratch Code Generation Tasks

the context of scratch tasks. Specifically, we annotated 41 scratch code generation tasks and provided clear descriptions for each in plain text. The remaining shapes were too complex to describe without ambiguity in plain text. Then, we compared the three modes of presenting the task, image only, text only, and the blend of an image and its description in text. Interestingly, for both GPT-4V and Gemini 1.5 Flash, the model performed worse when the task was presented only in the image, compared to the other modes. This phenomenon is counterintuitive as for humans, perceiving the images should be easier than first reading a description, imagining it, and then writing a code for it. Additionally, as presented in Table 4 the blend of image and text only slightly improved GPT-4V’s performance (from 38% to 40%). These two findings show that there is still much room for improvement especially in the visual components of LMMs.

B.4 Reasons of Failure

We manually investigated GPT-4V’s failures in solving Scratch tasks in a single run to find the major causes of failure. We find four major causes: 1) Shape identification error: where the model fails to completely capture existent shapes in the input image, for instance, if it confuses a semicircle with a circle or assigns non-existent shape attributes to the input image. 2) Counting error: where the model fails to count adequately, (e.g., three triangles counted as four), 3) Orientation error: where the model fails to correctly find the relationships between different components of a shape (e.g., semicircle on top of a square vs. at its bottom), and 4) Implementation error: where the model’s generated code does not follow the pre-planned pseudocode.

We manually investigated GPT-4V’s failure output in the scratch code generation task and the results are provided in Table 5, where the failures are not mutually exclusive as a model can perform a combination of errors in each task. Furthermore, while the first three errors are according to the vision component in these models, we see that 64% of the failures are according to these causes, and in 36% of failure cases, there are no apparent vision errors.

Cause	Description	Percentage
Shape identification error	Shape Identification Error: The model fails to completely capture existent shapes in the input image, confusing or misattributing shapes	25%
Counting error	Counting Error: The model inadequately counts the elements.	35%
Orientation error	Orientation Error: The model fails to correctly determine the spatial relationships between different components of a shape	21%
Implementation error	Implementation Error: The model’s generated code does not adhere to the pre-planned pseudocode, resulting in incorrect implementation.	45%

Table 5: Major Causes of GPT-4V’s Failures in Scratch Tasks; note that the failures are not mutually exclusive, as a model can perform a combination of errors in each task

C Experiment Setup

C.1 Automatic Evaluation of Code Output

Evaluation of the output code by an AI model is performed automatically. First, the output of the AI model is processed to extract the code piece of output. Then, this piece of code is run in a sandbox, and the shape produced by the code is stored. An illustration of this pipeline is provided in Figure 9. Finally, using the *OpenCV* module in Python, the binary versions of the correct shape and the produced shape are compared using an adjusted measure of bitwise similarity where we first use the bounding box technique with *OpenCV* to find the exact location of the shape and then calculate similarity with the formula:

$$\frac{|B_a \cap B_m|}{|B_a \cup B_m|}$$

where B_a and B_m represent black pixels in the input and LMM output, respectively. This metric measures the ratio of co-occurring black pixels to the total black pixels. Here, we utilize a heuristic approach in labeling the correctness of the model’s output. If the bitwise similarity between output and ground truth is higher than 95% the models’ output is labeled as correct and incorrect otherwise. To make sure that our heuristic in labeling the correctness of generated shapes is reliable, we manually annotated 2000 pairs of input and output images and we found that only three instances of pairs were labeled incorrectly (two of them false negative and the other false positive.), leading to an error rate of 0.15% which shows the high level of reliability in the heuristic we used.

C.2 Prompting

C.2.1 Basic Prompt

```
In each task, the user provides an image of an abstract geometric shape or pattern and an instruction, you need to generate a code in Python Turtle that follows the user's request.
```

Figure 2: basic prompt used in our experiments

C.2.2 v-CoT Prompt

```
You are Turtle Geometrician, you are an expert in reasoning about images and generating code in Python Turtle using images You need to follow the steps below before generating the answer:  
(1) Describe the relevant information from the image needed to answer the question. List all relevant artifacts from the image.  
(2) Use the information described in (1) to reason about the problem by working step by step to arrive at the final piece of code.  
(3) Generate the final code. NEVER use "pensize" function in your code.
```


Figure 3: v-CoT prompt used in our experiments

C.2.3 A Complete Example

Here we provide an instance of a complete prompt we used for a `tweak code generation` task with CoT prompting:

C.2.4 Arbitrary Output

Here we provide the CoT prompt we used for the model to provide a code in any arbitrary language or library that creates the desired shape.



```

System: You are Turtle Geometrician, you are an expert in reasoning about
images and generating code in Python Turtle using images. You need to
follow the steps below before generating the answer:
(1) Describe the relevant information from the image needed to answer the
question. List all relevant artifacts from the image.
(2) Use the information described in (1) to reason about the problem by
working step by step to arrive at the final piece of code.
(3) Generate the final code. NEVER use "pensize" function in your code.

Text: Provide a code in Python turtle that in the given shape inserts a
circle of an equal size to the smaller circle on the left of the bigger
circle to make a vertically symmetrical shape.

Complete the code:

import turtle
from math import *
t = turtle.Turtle()
large_circle_radius=100
small_circle_radius=50
...

```

Figure 4: An example of a complete prompt for a tweak code generation task with using v-CoT prompting.

```

You are an expert in reasoning about images and generating code in any language you
prefer. You need to follow the steps below before generating the code that answers
the user's request:
(1) Describe the relevant information from the image needed to answer the question.
List all relevant information from the image.
(2) Use the information described in (1) to reason about the problem by working
step by step to arrive at the final piece of code.
(3) Generate the final code. Your code can be in any visual language or library,
such as Matplotlib, TikZ, etc.

```

Figure 5: The system prompt we used for the results discussed in Section B.2

C.3 Rabbit

C.3.1 Prompt used

The prompt we used for this experiment is provided in Figure 6.

C.3.2 Definition of the class

The rabbit class is an arbitrary class that we defined based on *Turtle* class in the Python Turtle Module. This minimal set of functions includes all functions that a programmer or a model needs to create all of the tasks in TurtleBench. We defined this new set of functions to measure how GPT-4V is able to generalize its abilities in generating code in Python Turtle to a similar but minimally different set of functions.

```

import turtle

class Rabbit(turtle.Turtle):
    def __init__(self):
        super().__init__()
        self.setheading(90)
        self.pensize(5)
        self.hideturtle()

    def aa(self, length):

```

Suppose that I have a library named Rabbit in Python. Rabbit library has an object constructor named Rabbit which is an object that moves on the screen and draws lines. It only has these functions:

aa(length): goes front or back (if the length is negative) and draws a line with the length of pixels.

bb(degree): The rabbit turns its head right or left (if degree is negative).

cc(radius, degree): creates an arc with the given radius for the given degree. If degree=360 it creates a circle. The center of the circle is in the left of the rabbit.

pp(vanish): if vanish=True vanishes Rabbit object so if it moves does not draw anything, and if vanish=False, it appears the Rabbit object so if it moves draws on the screen.

you call the functions on an object of Rabbit, such as r.aa(length) where r is an object of Rabbit. When r is created, it faces north (up) on the screen and it does not vanish, so it is in drawing mode.

You are Rabbit Geometrician, you are an expert in reasoning about images and generating code in Python Rabbit using images. You need to follow the steps below before generating the answer:

- (1) Describe the relevant information from the image needed to answer the question. List all relevant artifacts from the image.
- (2) Use the information described in (1) to reason about the problem by working step by step to arrive at the final piece of code.
- (3) Generate the final code. Only use commands in the Rabbit class.

Figure 6: v-CoT prompt used for generalization experiments discussed in Section B.1

```

        self.forward(length)

def bb(self, degree):
    self.right(degree)

def cc(self, radius, degree):
    self.circle(radius, degree)

def pp(self, vanish):
    if vanish:
        self.penup()
    else:
        self.pendown()

```

C.4 Types of Tweak Tasks

TurtleBench includes a total of 130 tweak tasks. We provide a categorization for the tweaks as follows: There are five major types of tweaks in TurtleBench;

- Deletion: Removing a specified part of a shape
- Insertion: Adding a specific shape to the pattern as directed
- Rotation: Rotating the entire shape
- Reflection: Reflecting the entire shape or parts of it across specified lines
- Generalization: maintaining a pattern in the image constant while varying its parameters.

An illustration of instances of each type is provided in Figure 7. These types are not mutually exclusive as 10% of the tasks involve a combination of two types (e.g., removing one side of a square and inserting a semicircle instead). To successfully complete deletion and insertion tweaks, a model needs to demonstrate a nuanced understanding of the details in the image and program the resulting shape accordingly. In contrast, rotation tasks can be relatively easy as most of them can be solved only using a simple function in Turtle that can rotate the starting heading of the turtle which results in complete rotation in the entire shape (i.e., `turtle.right(angle)`).




Figure 7: Types of tweaks and their share in TurtleBench

C.5 Evaluating Image Complexity Using Contour Counts

As our result suggests that the vision component is contributing poorly to the models’ performance, to gain a better understanding of the visual obstacles for the models to solve the tasks, we defined a measure as a proxy for the complexity of shapes. For each provided image, we calculated the number of contours in each shape. In OpenCV, a contour is a curve joining all the continuous points (along the boundary), having the same color or intensity. Contours are a useful tool for shape analysis and object detection and recognition. The high number of contours in an image hints that there are many shapes being involved and interleaving with each other, which makes understanding and extracting underlying patterns challenging.

We calculated the number of contours in each shape by utilizing the corresponding function in OpenCV, and defined three arbitrary levels of complexity in the images, where the images which include only one contour (e.g., the basic square in Figure 1) are at level 1 (simple), images including less than 6 contours and more than 1 are at level 2 (medium) (e.g, the base shape of insertion example in Figure 7) and the images in which there are more than 6 contours (e.g., the base shape in generalization example in Figure 7) are at level 3 of the complexity (Complex). In Turtle, the proportions of complexity levels 1, 2, and 3 are 25%, 40%, and 35%, respectively.

We investigate how models perform over tweak tasks. There are 9 different ways that a pair of input and output image can combine. As shown in Table 6, the majority of tweak tasks (74) have same levels of complexity for the input and output image.

To examine how complexity of input and output shapes impact the results, we categorize tweak tasks in the 9 different categories and count the number of tasks that are ever solved by GPT-4V under any prompting method in code generation and code edit tasks during 6 different runs. As shown in Table 6, the more complex the input shape is, the more challenging solving the task is.

		Output Complexity		
		Simple	Medium	Complex
Input Complexity	Simple	35% (7/20)	30% (3/10)	25% (1/4)
	Medium	40% (2/5)	18% (6/33)	7% (1/13)
	Complex	20% (1/5)	11% (2/19)	19% (4/21)

Table 6: The number of tweak tasks under each category and the percentage of those tasks ever solved by GPT-4V in different settings.

C.6 Task Instances

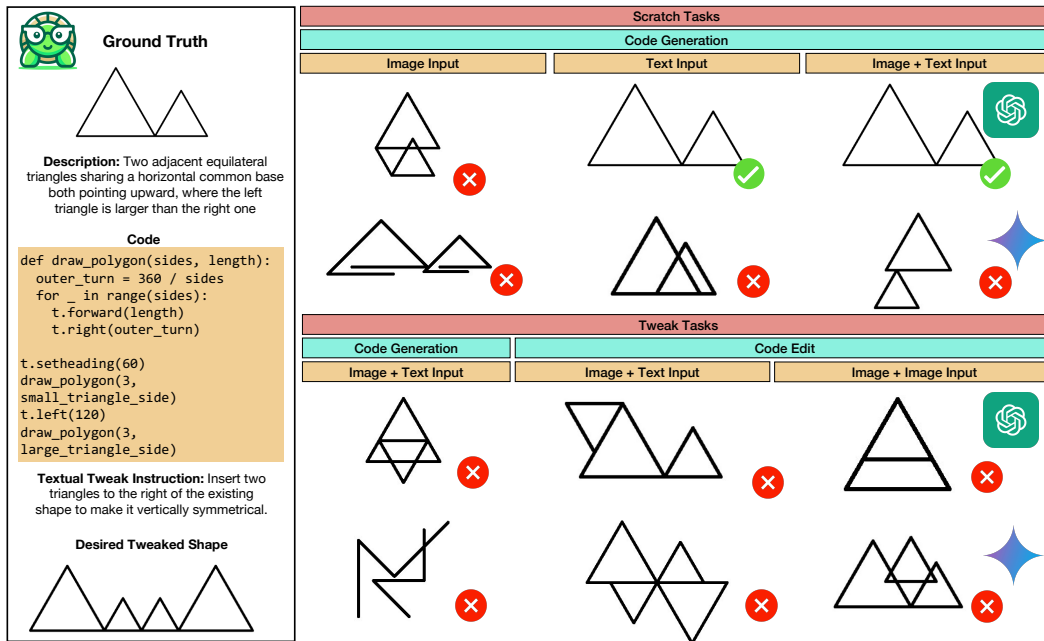


Figure 8: An illustration of different modes of a single task in TurtleBench along with the images generated by code from the outputs of GPT-4V and Gemini 1.5 Flash. More examples are provided in Appendix Figure 10

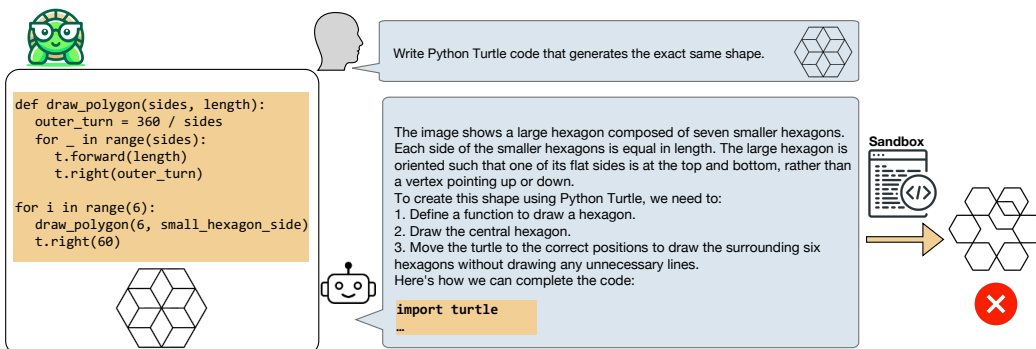





















































Figure 9: An illustration of our evaluation pipeline

 <p>Ground Truth</p>  <p>Description: A square with two horizontal and vertical sides, inscribed in a circle.</p> <p>Code</p> <pre>def draw_polygon(sides, length): outer_turn = 360 / sides for _ in range(sides): t.forward(length) t.right(outer_turn) draw_polygon(4, square_side) t.left(-135) t.circle(square_side/2*sqrt(2),-360)</pre> <p>Textual Tweak Instruction: Write a code in Python turtle that creates the given shape without the quarter circles on the left and the right of the square.</p> <p>Desired Tweaked Shape</p> 	<p>Scratch Tasks</p> <p>Code Generation</p> <table border="1"> <tr> <td>Image Input</td> <td>Text Input</td> <td>Image + Text Input</td> </tr> <tr> <td> ❌</td> <td> ❌</td> <td> ✅ </td> </tr> <tr> <td> ❌</td> <td> ✅</td> <td> ❌ </td> </tr> </table>			Image Input	Text Input	Image + Text Input	 ❌	 ❌	 ✅ 	 ❌	 ✅	 ❌ 		
	Image Input	Text Input	Image + Text Input											
 ❌	 ❌	 ✅ 												
 ❌	 ✅	 ❌ 												
<p>Tweak Tasks</p> <table border="1"> <tr> <td>Code Generation</td> <td colspan="2">Code Edit</td> </tr> <tr> <td>Image + Text Input</td> <td>Image + Text Input</td> <td>Image + Image Input</td> </tr> <tr> <td> ❌</td> <td> ❌</td> <td> ❌ </td> </tr> <tr> <td> ❌</td> <td> ❌</td> <td> ❌ </td> </tr> </table>			Code Generation	Code Edit		Image + Text Input	Image + Text Input	Image + Image Input	 ❌	 ❌	 ❌ 	 ❌	 ❌	 ❌ 
Code Generation	Code Edit													
Image + Text Input	Image + Text Input	Image + Image Input												
 ❌	 ❌	 ❌ 												
 ❌	 ❌	 ❌ 												




















































 <p>Ground Truth</p>  <p>Description: Four adjacent squares of an equal size that form a larger square.</p> <p>Code</p> <pre>def draw_polygon(sides, length): outer_turn = 360 / sides for _ in range(sides): t.forward(length) t.right(outer_turn) x = 4 for i in range(x): draw_polygon(4, small_square_side) t.right(360 / x)</pre> <p>Textual Tweak Instruction: Write a code in Python Turtle that creates the given shape without the small square on the top right.</p> <p>Desired Tweaked Shape</p> 	<p>Scratch Tasks</p> <p>Code Generation</p> <table border="1"> <tr> <td>Image Input</td> <td>Text Input</td> <td>Image + Text Input</td> </tr> <tr> <td> ✅</td> <td> ✅</td> <td> ✅ </td> </tr> <tr> <td> ❌</td> <td> ❌</td> <td> ❌ </td> </tr> </table>			Image Input	Text Input	Image + Text Input	 ✅	 ✅	 ✅ 	 ❌	 ❌	 ❌ 		
	Image Input	Text Input	Image + Text Input											
 ✅	 ✅	 ✅ 												
 ❌	 ❌	 ❌ 												
<p>Tweak Tasks</p> <table border="1"> <tr> <td>Code Generation</td> <td colspan="2">Code Edit</td> </tr> <tr> <td>Image + Text Input</td> <td>Image + Text Input</td> <td>Image + Image Input</td> </tr> <tr> <td> ✅</td> <td> ❌</td> <td> ❌ </td> </tr> <tr> <td> ❌</td> <td> ✅</td> <td> ❌ </td> </tr> </table>			Code Generation	Code Edit		Image + Text Input	Image + Text Input	Image + Image Input	 ✅	 ❌	 ❌ 	 ❌	 ✅	 ❌ 
Code Generation	Code Edit													
Image + Text Input	Image + Text Input	Image + Image Input												
 ✅	 ❌	 ❌ 												
 ❌	 ✅	 ❌ 												

Figure 10: Two examples of tasks in TurtleBench across different modalities