

Beyond Visual Understanding Introducing PARROT-360V for Vision Language Model Benchmarking

Harsha Vardhan Khurdula Basem Rizk Indus Khaitan

Redblock AI

Abstract

Current benchmarks for evaluating **Vision Language Models (VLMs)** often fall short in thoroughly assessing these models' abilities to understand and process complex visual and textual content. They typically focus on simple tasks that do not require deep reasoning or the integration of multiple data modalities to solve an original problem. To address this gap, we introduce the **PARROT-360V Benchmark**, a novel and comprehensive benchmark featuring 2487 challenging visual puzzles designed to test VLMs on complex visual reasoning tasks. We evaluated leading models—**GPT-4o**, **Claude-3.5-Sonnet**, and **Gemini-1.5-Pro**—using PARROT-360V to assess their capabilities in combining visual clues with language skills to solve tasks in a manner akin to human problem-solving. Our findings reveal a notable performance gap: state-of-the-art models scored between **28% to 56%** on our benchmark, significantly lower than their performance on popular benchmarks. This underscores the limitations of current VLMs in handling complex, multi-step reasoning tasks and highlights the need for more robust evaluation frameworks to advance the field.

1 Introduction

Vision Language Models (VLMs) have shown remarkable capabilities in integrating visual and textual data, excelling in tasks like image captioning and object recognition (Wang et al., 2024).

The aspiration to create artificial intelligence that can seamlessly integrate into daily life—solving problems, performing tasks, and providing expert knowledge—has long been a driving force in technological advancement (Mintz and Brodie, 2019). Recent developments in VLMs have brought us closer to this vision, showcasing impressive abilities in understanding and generating both textual and visual data (Yang et al., 2024a). Models like GPT-4o, Claude-3.5-Sonnet, and Gemini-1.5-Pro

have set new standards in the field while showcasing high performance for vision-related benchmarks.

The rapid evolution of these models has sparked concerns. There is a growing fear that AI could replace human labor (Eloundou et al., 2024). These fears are often based on hypothetical scenarios rather than current capabilities. Despite these apprehensions, it's crucial to assess whether these models truly perform at critically claimed levels, especially in complex tasks that mirror real-world challenges.

Our benchmark, **PARROT-360V**, contributes to evaluating leading VLMs by focusing on step-by-step visual reasoning tasks. We aim to identify gaps between reported capabilities and actual performance, offering insights into specific areas where these models may underperform.

2 Why A New Benchmark?

Many commonly used benchmarks such as **MMMU** by Yue et al. (2024), **ChartQA** by Masry et al. (2022), and **AI2D** by Kembhavi et al. (2016) have been designed to evaluate VLMs on tasks that are limited in scope, such as basic image-text alignment or single-step reasoning. These benchmarks are typically straightforward, and models can at times overfit to the datasets, resulting in misleadingly high-performance scores (Samuel et al., 2024). We aim to adequately test models on puzzles that require the skills of image-text alignment, multi-step reasoning, and sequential logic handling. In particular order, they correspond to sub-tasks that are critical for real-world decision-making (Tu et al., 2024) and the analysis steps required.

2.1 Challenges In Reproducibility For VLMs

Reproducibility is a significant challenge in evaluating vision-based tasks, especially when dealing with VLMs (Yang et al., 2024b). Unlike purely

textual models, vision models rely on visual input, which can be subject to variability in data preprocessing, annotation, and context (Wu et al., 2024). This makes it difficult to replicate the exact conditions under which a model achieves specific results for other benchmarks.

A lack of standardization in how input images are processed or prompts are structured can lead to discrepancies in model outputs when evaluated across different platforms or test environments (Anagnostidis and Bulian, 2024). Moreover, existing benchmarks, which often focus on answering a question posed in the text within an image or on top of the content depicted in it (a graph), do not adequately capture the abilities of the model. The output for a question is heavily skewed by the size of the data it was trained on. Rather benchmarking for VLMs should evaluate perception, and ability to use that information. And move away from the emphasis of answering multiple choice questions from an image, which could have easily been represented as a question in text (Yue et al., 2024).

2.2 A Fairer Evaluation Paradigm

To ensure fairer comparisons, a benchmark should not just test how much knowledge a model has absorbed but should also evaluate how well it can perceive and follow instructions based on the visual inputs provided. This is where our PARROT-360V shines, as it requires models to integrate visual perception with textual reasoning, testing their ability to interpret, reason, and solve complex problems step by step, rather than regurgitating memorized knowledge (Duan et al., 2023).

This shift in focus is crucial for evaluating VLMs in a way that reflects their actual capabilities in real-world scenarios, where their performance must rely on accurate perception and decision-making rather than simply having been trained on vast quantities of data and determining its strengths towards being employed for automation involving visual tasks (Schwartz et al., 2023).


The **PARROT-360V** benchmark represents a step forward in addressing key challenges in reproducibility, data bias, and unfair comparisons in the field of VLMs. It provides a rigorous and fair benchmark for evaluating vision models based on their perceptual and reasoning abilities by employing Chain-of-Thought (CoT) (Wei et al., 2022) to plan how the model is going to solve the current puzzle, offering a clearer picture of how these mod-

JUMBLE FOR KIDS

THAT SCRAMBLED WORD GAME!®

By David L. Hoyt and Jeff Knurek

The letters of these crazy words are all mixed up. To play the game, put them back into the right order so that they make real words you can find in your dictionary. Write the letters of each real word under each crazy word, but only one letter to a square.



Now you're ready to solve today's Jumble For Kids. Study the picture for a hint. Play around with the letters in the circles. You'll find you can put them in order so that they make your funny answer.

Print your answer here: “ - ”

Figure 1: Sample from the PARROT-360V Dataset.

els would perform in real-world applications (Wei et al., 2023).

3 PARROT-360V Dataset

The **PARROT-360V Benchmark** dataset was carefully curated by scraping Jumble puzzles from the internet to challenge VLMs in solving complex jumbles (Redblock.ai, 2024). Each scraped puzzle combines various elements representing an instance of gameplay, as shown in Figure 3, which serves as the input puzzle to the VLM. Furthermore, we extract additional features from the solved puzzle to obtain the ground truth labels, as shown in Figure 4.

The dataset corresponds to an established collection of puzzles, all of which are of the same format, with each containing a different set of clues/jumbled words, and a bonus clue with a visual component (associated with a cartoon) as shown in Figure 1. The dataset is intended to evaluate not only language understanding but also visual perception and reasoning, making it a more rigorous test for these models.

Field	Description
Date	The date the puzzle was published.
PDF File Path	The location of the puzzle in PDF format.
Question Screenshot	A visual representation of the puzzle’s clues.
Answer Screenshot	A screenshot of the correct solution for comparison purposes.
Clues (Clue_1 to Clue_4)	The four scrambled words presented as clues in the puzzle.
Answers (Answer_1 to Answer_4)	The correct solutions to the jumbled clues.
Visual Clue	A scrambled phrase or word that is solved using characters from the answers.
Puzzle Answer	The correct solution to the bonus clue is derived from the characters circled in the answers.

Table 1: Structure of the PARROT-360V Dataset, detailing the fields contained in each puzzle entry.

3.1 Data Structure

Each puzzle in the PARROT-360V dataset consists of the following features:

- **Regular Clues:** These are scrambled words given as clues within the puzzle image, which the models must unscramble to find the correct words, and extract characters circled characters to form a bonus clue.
- **Visual Clue:** A cartoon or image that contains a visual hint. The models must interpret this image or other relevant information to form a bonus answer.
- **Answer Constraints:** The models are required to piece together specific letters (often circled in the image) from the unscrambled words to form the bonus answer.

The dataset contains 2487 samples, with 14 distinct features from the release date of the puzzle to the annotated answers as text within the dataset, as described in table 1.

4 PARROT-360V Benchmarking Setup

We applied our benchmark on three state-of-the-art VLMs—**GPT-4o**, **Claude-3.5-Sonnet**, and **Gemini-1.5-Pro**. Requiring the models to integrate both visual and textual information to arrive at correct answers, the evaluation environment is designed to simulate a realistic puzzle-solving environment:

- **Input:** Images of jumbled word puzzles, including visual clues (circled letters, and cartoon characters).

- **Task:** The models are required to solve the scrambled words, interpret the visual clues, and synthesize the final solution by forming bonus answers from the circled letters in the puzzle. While framing context from the posed question about the cartoon (as shown in the algorithm 1).
- **Metrics:** We measured the correctness of their responses and the proportion of hallucinations, where the model incorrectly used characters that were not part of the given clues. Additionally, we calculated each model’s overall performance across multiple dimensions—accuracy, sequential performance (evaluation of intermediate steps involved), and hallucination rate.

To mitigate the issue of potential data contamination, we ensured that the setup used was entirely novel and tested the models for real-world tasks, that cannot be reproduced by simply using its knowledge base. Rather we prompt the VLM to explicitly/step-by-step address the required tasks: identify characters, plan how it is going to handle the bonus clue, and solve the entire puzzle.

One of the core challenges presented by the PARROT-360V benchmark is for VLMs to identify the circled letters within the image clues as shown in Figure 1 for a model. These letters are crucial for solving the bonus clue, and failure to correctly identify these letters often results in hallucinations or incorrect answers from the models.

In addition to recognizing circled characters, models are expected to interpret visual information from the cartoon or accompanying image (Figure

Algorithm 1 TASK: Solve Jumble Puzzle

Input: Image containing jumbled words W_1, W_2, W_3, W_4 , circled letter positions, bonus section, and a visual clue.

Output: Unscrambled words U_1, U_2, U_3, U_4 , reasoning behind the bonus answer B

Planning Phase:

Recognize that each word needs to be unscrambled and circled letters extracted to form the bonus clue.

Execution Phase:

for $i = 1$ to 4 **do**

 Unscramble W_i to get U_i

 Extract the circled letter from U_i at position P_i

end for

Concatenate circled letters to form bonus clue C

Unscramble C to get the bonus answer B while considering the posed question and visual clue.

1). This requires not only extracting individual characters but also understanding the broader context of the image to form a coherent solution to the puzzle. These puzzles that require visual understanding, multi-step reasoning, and sequential logic challenge the perception that current VLMs excel in complex, real-world tasks. This plays a critical role in understanding the limitations of these models in tasks that go beyond simple image-text QA (Yue et al., 2024).

4.1 Metrics

To quantify the models' performance on PARROT-360V, we developed a scoring system that assigns **weights** to each component of the puzzle. The scoring system is designed to reflect the importance of each task and to penalize omissions or incorrect answers appropriately.

Scoring components each puzzle consists of:

- **Four Scrambled Words:** Each worth 10 points. (There are four scrambled words/clues, within each puzzle thus a candidate can score 40 points at the most in this section.)
- **Synthesizing Answers to Extract Key Characters:** Worth 10 points. (Each unscrambled clue serves as a distinct answer. Certain characters are circled within these answers; concatenating these circled characters provides the final bonus clue.)
- **Solution to the Puzzle:** Worth 20 points. (Using the extracted characters and interpreting the cartoon, VLMs gather and synthesize information to solve the puzzle.)

The total possible points for each puzzle are 70 points.

4.2 Scoring Methodology

For each VLM's attempt at solving a puzzle, we applied the following evaluation criteria:

- **Correct Answer:** If the model's answer matches the labeled correct answer exactly (case-insensitive), it receives full points for that component¹.
- **No Answer or Incorrect Answer:** If the model provides no answer or an incorrect answer, it receives a penalty of -5 points for that component.
- **Negative Score Adjustment:** If the total points earned for a puzzle are negative due to penalties, the score is clipped to zero.
- **Normalization:** The total points earned are divided by the total possible points (70) to obtain a performance score between 0 and 1, rounded to two decimal places.

Let:

- W_i be the weight for component i .
- A_i is the model's answer for component i .
- C_i be the correct answer for component i .
- T be the total possible points (70).

¹Clue or One of the sequential tasks involved in solving the given puzzle.

Model	MMMUS	Mathvista	AI2D	ChartQA	Average Performance
GPT-4o	0.69	0.64	0.94	0.86	0.78
Claude 3.5 Sonnet	0.68	0.68	0.95	0.91	0.80
Gemini 1.5 Pro	0.62	0.64	0.81	0.81	0.72

Table 2: Performance of different models across various benchmarks.

For each component:

$$P_i = \begin{cases} W_i & \text{if } A_i \text{ is exactly } C_i, \\ -5 & \text{if } A_i \text{ is incorrect or missing.} \end{cases}$$

Total Points Earned:

$$P_{\text{total}} = \sum_i P_i \quad (1)$$

Clipped Total Points Earned:

$$P_{\text{adjusted}} = \max(0, P_{\text{total}}) \quad (2)$$

Performance Score:

$$PARROT360V_{\text{Score}} = \frac{P_{\text{adjusted}}}{T} \quad (3)$$

Hallucination within PARROT-360V benchmarking is related to the frequency with which models introduced information not present in the input. The hallucination rate is the error rate, i.e. the proportion of incorrect predictions given by a VLM:

$$HallucinationRate = 1 - (PARROT360V_{\text{Score}}) \quad (4)$$

5 Results

The evaluation of GPT-4o, Claude-3.5-Sonnet, and Gemini-1.5-Pro on our benchmark highlights the significant limitations of existing benchmarks in capturing true multimodal reasoning abilities. Unlike tasks found in common benchmarks like **MMMUS**, **MathVista**, or **ChartQA** (as shown in the table 2), PARROT-360V places special emphasis on complex, multi-step reasoning involving visual puzzles. This difference is reflected in the sharp decline in performance when these models are tested on our benchmarking dataset.

On benchmarks such as **MMMUS** and **MathVista**, GPT-4o and Claude-3.5-Sonnet achieved high scores of 0.69 and 0.72, respectively (Table 2), mainly because these tasks focus on simple image-text alignment or basic reasoning. In these tests,

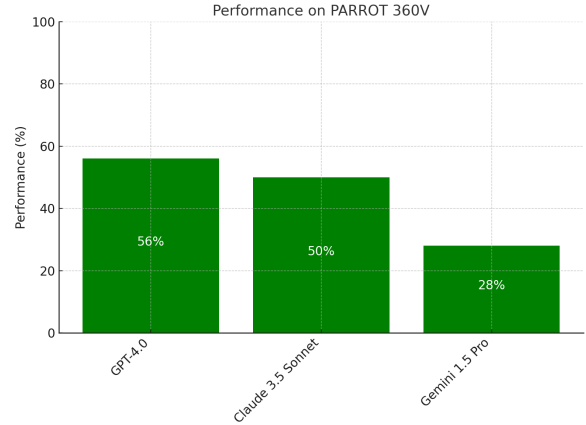


Figure 2: Performance of State-of-The-Art VLMs on PARROT-360V

the image often serves merely as a backdrop to a question that could just as easily be presented as pure text, reducing the need for genuine visual understanding.

PARROT-360V, by contrast, involves complex tasks like word unscrambling, bonus clue extraction, and interpreting visual elements, all requiring deep integration of visual and textual information. GPT-4o’s accuracy dropped significantly to **0.57**, Claude-3.5-Sonnet’s to **0.50**, and Gemini-1.5-Pro’s to **0.28** (Figure 2) in this more challenging setup, demonstrating how existing benchmarks fail to reflect the complexity required for real-world tasks.

5.1 Visual Perception Failures

In tasks such as identifying circled letters within the puzzle which is an image (Figure 1), all three candidate models struggled. For instance, Gemini-1.5-Pro exhibited a high hallucination rate of **72%**, largely due to its inability to accurately recognize and use visual inputs. This stands in contrast to simpler benchmarks like **AI2D**, where models face straightforward visual questions with limited need for complex image interpretation (Kembhavi et al., 2016). Often the models also failed to synthesize visual scrambled text effectively. The challenge of extracting circled letters and forming a correct bonus clue required higher-order detail in reason-

ing, which is not typically tested in conventional benchmarks.

5.2 Hallucination Issues

All three VLMs exhibited frequent hallucinations during evaluation, particularly when trying to derive answers from visual cues. While tasks in **ChartQA** or **MathVista** are often solved by applying memorized data or pattern recognition, our benchmark exposed the models' limitations in handling dynamic, real-time visual information. GPT-4o had a hallucination rate of **43%**, Claude-3.5-Sonnet with **50%**, and Gemini-1.5-Pro with **72%** as shown in the figure 2, thus proving how heavily models depend on structured data rather than real reasoning from raw inputs.

6 Discussion

When we compared the results of our proposed benchmark to existing benchmarks like MMMU, ChartQA, MathVista, and AI2D, it became clear that these benchmarks don't truly test a model's ability to reason through complex, real-world visual problems. In benchmarks like MMMU, models often achieve high scores (e.g., GPT-4o scoring 0.69 as shown in table 2) because the tasks typically involve static image-text alignment or basic pattern recognition. These benchmarks don't require the models to think but rather to retrieve memorized information or recognize patterns from training data. In essence, they reduce the challenge to answering questions that could just as easily be text-based.

In contrast, our framework challenges the models with multi-step reasoning and visual puzzles that demand a higher level of understanding. Models can't just rely on large datasets or pre-learned patterns; they need to synthesize information from both text and images. For example, tasks like identifying circled letters in images and using them to solve a bonus clue are far more reflective of real-world complexity than simple image-caption matching. When we saw models like Gemini-1.5-Pro struggle with hallucinations (72% rate as seen in figure 2), it was a clear indication that they're not truly equipped for these kinds of tasks—yet these are the tasks that matter when it comes to applying AI in fields like healthcare or automation.

One of the biggest takeaways from PARROT-360V was that the models performed significantly worse on our benchmark—with performance scores as low as **28%** for Gemini-1.5-

Pro—compared to traditional benchmarks (Figure 2). In short, current benchmarks are giving us an incomplete and often inflated view of what these models can do. The models' performance drop on PARROT-360V proves that while they might excel at answering questions from pre-learned data, they struggle when it comes to reasoning through complex, multi-step visual tasks. To move forward, we need benchmarks like our PARROT-360V that challenge these models to think and reason, not just recognize or recall.

7 Conclusion

With **PARROT-360V**, we aim to push the boundaries of how we evaluate VLMs by focusing on real-world tasks that demand visual perception, multi-step reasoning, and instruction-following. We saw that traditional benchmarks are not enough. They tend to focus on simpler tasks like image-text alignment and QA, which doesn't truly challenge the models' ability to understand and process both visual and textual data together. In contrast, PARROT-360V makes models tackle tasks that require actual reasoning and visual integration, such as solving word puzzles with visual clues.

Our findings reveal that GPT-4o (56%), Claude-3.5-Sonnet (50%), and Gemini-1.5-Pro (28%) struggle on our benchmark when handling complex, real-world tasks. This performance gap underscores the need for a more reliable benchmark, which PARROT-360V aims to provide.

8 Limitations

While **PARROT-360V** introduces a fresh approach to evaluating VLMs, we recognize that there are areas where further refinement can enhance its robustness. One aspect we've observed is the task complexity. Puzzles within the proposed benchmarking dataset are intentionally challenging, and designed to test multi-step reasoning and visual perception. However, there are instances where the complexity may obscure whether a model's failure is due to genuine reasoning difficulties or simply the task's intricacy. As we move forward, maintaining the right balance in task difficulty will help ensure we accurately measure a model's reasoning capabilities.

Another focus is visual perception, as models must interpret visual clues and recognize circled letters. We aim to separate perception from reasoning to ensure fair evaluation. Lastly, to address

data contamination, we will regularly update the benchmark with new tasks to test models on unseen data.

Acknowledgments

We would like to acknowledge Aviral Srivastava, Raj Khaitan, and Jay Khatri from our team at Redblock along with Janit Anjaria, our AI advisor, for their valuable contributions to PARROT-360V. Their insights and support were instrumental in advancing our research.

References

- Sotiris Anagnostidis and Jannis Bulian. 2024. [How susceptible are llms to influence in prompts?](#) *Preprint*, arXiv:2408.11865.
- Haonan Duan, Adam Dziedzić, Nicolas Papernot, and Franziska Boenisch. 2023. [Flocks of stochastic parrots: Differentially private prompt learning for large language models](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 76852–76871. Curran Associates, Inc.
- Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. 2024. [Gpts are gpts: Labor market impact potential of llms](#). *Science*, 384(6702):1306–1308.
- Aniruddha Kembhavi, Mike Salvato, Eric Kolve, Minjoon Seo, Hannaneh Hajishirzi, and Ali Farhadi. 2016. [A diagram is worth a dozen images](#). *Preprint*, arXiv:1603.07396.
- Ahmed Masry, Do Xuan Long, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. 2022. [Chartqa: A benchmark for question answering about charts with visual and logical reasoning](#). *Preprint*, arXiv:2203.10244.
- Yoav Mintz and Ronit Brodie. 2019. [Introduction to artificial intelligence in medicine](#). *Minimally Invasive Therapy & Allied Technologies*, 28(2):73–81. PMID: 30810430.
- Redblock.ai. 2024. Parrot-360v benchmark. Available at: <https://huggingface.co/datasets/RedBlock/parrot360v>.
- Vinay Samuel, Yue Zhou, and Henry Peng Zou. 2024. [Towards data contamination detection for modern large language models: Limitations, inconsistencies, and oracle challenges](#). *Preprint*, arXiv:2409.09927.
- Sivan Schwartz, Avi Yaeli, and Segev Shlomov. 2023. [Enhancing trust in llm-based ai automation agents: New considerations and future challenges](#). *Preprint*, arXiv:2308.05391.
- Yahan Tu, Rui Hu, and Jitao Sang. 2024. [Ode: Open-set evaluation of hallucinations in multimodal large language models](#). *Preprint*, arXiv:2409.09318.
- Yiqi Wang, Wentao Chen, Xiaotian Han, Xudong Lin, Haiteng Zhao, Yongfei Liu, Bohan Zhai, Jianbo Yuan, Quanzeng You, and Hongxia Yang. 2024. [Exploring the reasoning abilities of multimodal large language models \(mllms\): A comprehensive survey on emerging trends in multimodal reasoning](#). *Preprint*, arXiv:2401.06805.
- 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 its reasoning in large language models](#). *Preprint*, arXiv:2201.11903.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). *Advances in neural information processing systems*, 35:24824–24837.
- Xiyang Wu, Souradip Chakraborty, Ruiqi Xian, Jing Liang, Tianrui Guan, Fuxiao Liu, Brian M. Sadler, Dinesh Manocha, and Amrit Singh Bedi. 2024. [Highlighting the safety concerns of deploying llms/vlms in robotics](#). *Preprint*, arXiv:2402.10340.
- Enneng Yang, Li Shen, Guibing Guo, Xingwei Wang, Xiaochun Cao, Jie Zhang, and Dacheng Tao. 2024a. [Model merging in llms, mllms, and beyond: Methods, theories, applications and opportunities](#). *Preprint*, arXiv:2408.07666.
- Qian Yang, Weixiang Yan, and Aishwarya Agrawal. 2024b. [Decompose and compare consistency: Measuring vlms’ answer reliability via task-decomposition consistency comparison](#). *Preprint*, arXiv:2407.07840.
- 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 Wenhao Chen. 2024. [Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi](#). *Preprint*, arXiv:2311.16502.

A Scraping And Puzzle Curation

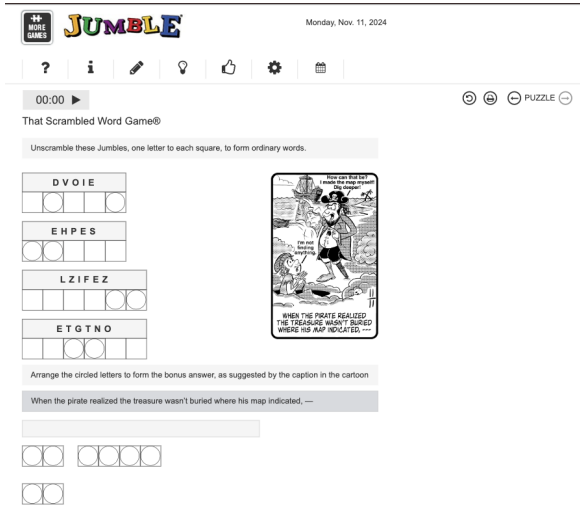


Figure 3: Snapshot of the Puzzle.

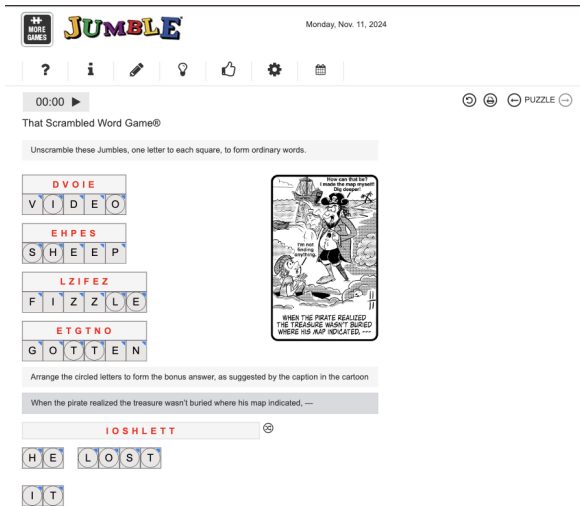


Figure 4: Snapshot of the Solved puzzle.