Visual Graph Arena: Evaluating Visual Conceptualization of Vision and Multimodal Large Language Models

Zahra Babaiee^{*1} Peyman M. Kiasari^{*1} Daniela Rus² Radu Grosu¹

Abstract

Recent advancements in multimodal large language models have driven breakthroughs in visual question answering. Yet, a critical gap persists, 'conceptualization'-the ability to recognize and reason about the same concept despite variations in visual form, a basic ability of human reasoning. To address this challenge, we introduce the Visual Graph Arena (VGA), a dataset featuring six graph-based tasks designed to evaluate and improve AI systems' capacity for visual abstraction. VGA uses diverse graph layouts (e.g., Kamada-Kawai vs. planar) to test reasoning independent of visual form. Experiments with state-of-theart vision models and multimodal LLMs reveal a striking divide: humans achieved near-perfect accuracy across tasks, while models totally failed on isomorphism detection and showed limited success in path/cycle tasks. We further identify behavioral anomalies suggesting pseudo-intelligent pattern matching rather than genuine understanding. These findings underscore fundamental limitations in current AI models for visual understanding. By isolating the challenge of representationinvariant reasoning, the VGA provides a framework to drive progress toward human-like conceptualization in AI visual models. The Visual Graph Arena is available at: vga.csail.mit.edu.

1. Introduction

In recent years, the fields of object recognition (Girshick et al., 2015) and visual question answering (Antol et al., 2015; Goyal et al., 2017; Yang et al., 2022) have witnessed remarkable progress, pushing the boundaries of artificial intelligence. Despite these advancements, a gap remains between current AI capabilities and the goal of achieving artificial general intelligence (AGI). One critical aspect of this gap is the ability of AI models to engage in deep reasoning and understand fundamental concepts from visual inputs (Gupta & Kembhavi, 2023), which is essential for reaching human-like cognitive abilities (Han et al., 2019). This gap is particularly evident in what we term 'conceptualization' - the ability to recognize and reason about concepts across different representations, even when their visual form varies significantly.

Humans have a long history of developing various methods of visual representation to simplify complex information and enhance understanding. These representations span across diverse domains, such as chemical structures, logic circuits, flowcharts, and statistical data, each with its unique set of visual conventions and abstractions. Figure 1 showcases the diversity and complexity of visual representations that humans regularly encounter and interpret. In Figure 1-a. two different visualizations of the same logic circuit are presented, illustrating how identical functional relationships can be conveyed through varied graphical styles. Similarly, Figure 1-b depicts two representations of the same chemical structure, highlighting the challenge of recognizing underlying molecular configurations when presented in different visual forms. Figure 1-c features a graph drawn in two distinct layouts, both representing the same underlying data structure and relationships. These examples underscore a crucial aspect of visual reasoning and highlight the core challenge of 'conceptualization': the ability to consistently understand different visual representations of the same concept, a task that is typically intuitive for humans but remains challenging for AI.

Defining conceptualization: Undoubtedly, the ability to generalize beyond superficial patterns lies at the heart of human-like reasoning. While out-of-distribution generalization encompasses various challenges, we focus specifically on models' ability to understand different representations of the same underlying concept. To precisely characterize this aspect, we use the term "Conceptualization".

Conceptualization is a specific subset of out-of-distribution (OOD) generalization that examines a model's ability to recognize and reason about OOD data whose *underlying*

^{*}Equal contribution ¹CPS, TU Wien, Vienna, Austria ²CSAIL, MIT. Correspondence to: Peyman Kisari <peyman.kiasari@tuwien.ac.at>.

Proceedings of the 42^{nd} International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).







(a) Two representations of a logic circuit. (b) Two depictions of a chemical structure. (c) A graph drawn in two different layouts.

Figure 1: Examples of different visual representations of identical concepts across different scientific and technical domains. Each pair of images illustrates the challenge of recognizing identical relationships or structures when presented in varied visual formats.

concept aligns with the training distribution, but whose representation diverges in form, structure, or manifestation.

Unlike broader OOD generalization challenges—which may involve shifts in spurious correlations, contextual biases, or surface-level features—conceptualization isolates the capacity to abstract and transfer core conceptual invariants across representational variations. (Figure 1)

For instance, while understanding different accents in speech recognition represents an OOD challenge, it falls outside the scope of conceptualization, as the core linguistic constructs remain consistent. Indeed, we would not attribute an American's difficulty in understanding a Scottish accent to a lack of conceptual understanding of English.

Why Graphs? Among the various visual representations, graphs, consisting of nodes and edges, offer an ideal test case for studying conceptualization in AI systems. While these graphical representations are easily interpreted by humans, particularly when the graphs are simple and have a limited number of nodes, they present a challenge for AI models - the ability to abstract the underlying structure regardless of visual layout.

It is worth noting that while computational algorithms exist for solving graph problems through matrix operations or adjacency list traversal, our focus here is fundamentally different. The challenge lies not in developing graph-theoretic solvers, we aim to evaluate whether AI models can solve these problems using only visual reasoning and conceptual abstraction, without relying on algorithmic computation. This distinction is crucial to align more closely with humanlike reasoning.

This insight led to the creation of the "Visual Graph Arena", a dataset and benchmark consisting of six graph-based tasks. These tasks are designed to evaluate and enhance the visual reasoning and conceptualization capabilities of AI models through various challenges. The datasets contain a diverse range of training samples, from 27,000 to 150,000, ensuring robust training environments. The graphs have 8-9 nodes, small enough such that is easy to visually inspect their properties, but still large enough to create enough samples for the datasets.

The tasks in the Visual Graph Arena revolve around three main concepts: isomorphism, path, and cycle, with two subtasks for each concept. The training sets are drawn using specific layouts (Kamada-Kawai or random, depending on the task), while the test sets employ different layouts to assess whether models can learn and apply the concepts, regardless of their visual presentation, similarly to the human capabilities.

We evaluated several leading vision models, including ViT (Dosovitskiy et al., 2020), Swin Transformers (Liu et al., 2021), and ConvNeXt (Liu et al., 2022), on the Visual Graph Arena datasets. The results were very revealing: All models failed on the graph isomorphism detection tasks. Additionally, these models underperformed relative to the human capabilities on the other tasks, when tested with graphs drawn differently than the layouts in their training sets. Our investigation also revealed a fundamental reasoning gap in current state-of-the-art vision and multimodal language models. When we tested advanced AI models, such as GPT-o1 of OpenAI and Claude-3.5-Sonnet of Anthropic, on tasks that humans handle intuitively based on visual information alone, the results were striking. In a graph isomorphism detection task, where pairs of graphs were drawn in different layouts (Kamada-Kawai (Kamada et al., 1989) and planar), all tested models failed to recognize the identical underlying structures, while human subjects achieved more than 90% accuracy. This stark contrast between human and AI performance highlights a critical challenge in AI conceptualization.

The Visual Graph Arena dataset helps identify AI systems' limitations in processing conceptual visual data while providing a foundation for improvement. By testing models' conceptual reasoning abilities, we work to bridge the gap between human and machine interpretation of abstract graphics.



Figure 2: Illustration of a single graph displayed in three distinct visualization layouts to highlight the impact of graphical representation on perceived structure and complexity.

2. Related Work

Graph Datasets. Graph datasets are a cornerstone in numerous research domains (Hu et al., 2020), ranging from social network analysis to protein micro molecules (Morris et al., 2020). These datasets are typically divided into two main types: those that represent real-world information in graph form, like social media connections or road networks, and synthetic graph datasets crafted for specific computational experiments.

Among synthetic datasets, the variety is broad. For instance, various works have introduced datasets aimed at tasks like counting substructures within graphs (Chen et al., 2020) or distinguishing between different families of non-isomorphic graphs (Wang & Zhang, 2023; Abboud et al., 2020; Balcilar et al., 2021; Murphy et al., 2019). These synthetic datasets are primarily designed to benchmark and improve graph neural networks (GNNs). They often include data of nodes, edges, and additional features to facilitate computational processing and analysis tailored to GNN architectures.

However, a gap exists in these offerings concerning visual interpretation. Most existing graph datasets are structured for algorithmic or GNN consumption and do not address the challenge of visual graph reasoning directly. They are optimized for models that process graphs as textual or abstract data structures rather than as visual entities.

Visual Reasoning Datasets. Recent advancements in machine learning have seen a burgeoning interest in tasks that interrogate visual scenes through question-answering (QA) formats (Johnson et al., 2017; Antol et al., 2015).

Visual QA tasks have driven progress in multimodal reasoning, with datasets like CLEVR (spatial/relational reasoning) and FigureQA (scientific plots) using synthetic data to isolate challenges (Johnson et al., 2017; Kahou et al., 2017).

Recent benchmarks address diverse aspects: long-context document understanding (MMLONGBENCH-DOC), chain-

of-thought reasoning (Visual CoT), compositional robotic reasoning (ClevrSkills), mathematical Olympiad problems (Children's Math Olympiads evaluation), commonsense riddles (Visual Riddles), and multimodal math reasoning (MATH-Vision) (Gong et al., 2024; Shao et al., 2024; Guetta et al., 2024; Wang et al., 2024; Ma et al., 2024; Cherian et al., 2024; Haresh et al., 2024). A recent study introduces Vision Graph, a benchmark for visual reasoning in graph theory problems (Li et al., 2024), with two key limitations: its circular graph layouts become visually ambiguous with increasing nodes, and its dataset size is insufficient for training data-intensive models.

Our work introduces the "Visual Graph Arena", a graphbased dataset designed not only to evaluate AI model capabilities in solving graph theory problems but also, at a higher level, to assess their conceptualization ability. This is achieved by systematically varying graph layouts during both training and testing, thereby avoiding overfitting to specific patterns and isolating reasoning challenges from perceptual artifacts. By emphasizing scalability (27,500–155,718 samples per task) and layout diversity, we ensure that all samples remain visually solvable while probing deeper cognitive processes. This focus on practical solvability—validated through human evaluation—distinguishes our benchmark as a robust tool for advancing visual graph understanding and fostering models capable of generalizing across different visual contexts.

3. Visual Graph Arena

This section introduces the Visual Graph Arena Dataset, meticulously crafted to challenge and evaluate the visual reasoning capabilities of AI models through a series of graph-based tasks. The dataset is structured around three primary concepts, each divided into two tasks, resulting in a total of six distinct datasets. These datasets are designed to rigorously test various aspects of graph understanding and



Figure 3: Overview of the 'Visual Graph Arena Benchmark' detailing the division of tasks into categories of Graph Isomorphism, Graph Path Finding, and Graph Cycle Finding. Each category further breaks down into two subtasks.

reasoning.

In order to ensure that the graphs are both manageable and visually interpretable, we have constrained the number of nodes to range between 8 and 9. This node count strikes a balance between maintaining visual simplicity and providing sufficient complexity for generating robust training and test sets, which are suitable for deep learning applications.

The tasks are devised around core graph concepts that are intuitively recognizable by humans when interacting with small graphs:

- 1. **Graph Isomorphism:** Tests the ability to determine the isomorphic nature of graph pairs.
- 2. **Path Finding:** Focuses on the capability to identify various path types within the graphs.
- 3. Cycle Finding: Challenges the models to detect different types of cycles.

Figure 3 shows an overview of the Visual-Graph-Arena Benchmark tasks and datasets, and Table 1 shows the breakdown of the tasks with the number of samples in the training and test set for each task.

In the subsequent subsections, we will describe each task in detail.

3.1. Graph-Isomorphism Tasks

This task focuses on graph isomorphism, utilizing graphs consisting of 8-9 nodes, to test the models' capabilities in recognizing isomorphic structures. Graph isomorphism refers to the concept where two distinct graph representations share the same number of vertices, edges, and connectivity patterns, essentially making them indistinguishable in graph-theoretic terms despite their different visual layouts. To assess this, the task presents pairs of graphs drawn in different layouts, one in Kamada-Kawai and the other in planar, and models must determine whether the graphs are isomorphic.

The task is structured to ensure a balanced approach: each set contains an equal number of isomorphic and nonisomorphic graph pairs. This balance is crucial to prevent models from learning and exploiting label distribution biases instead of mastering the intended isomorphic detection.

Difficulty Levels: The graph-isomorphism-recognition task demands quite sophisticated reasoning from the models, as superficial attributes like edge count or node degree are insufficient for determining isomorphism. The models must engage in deeper analytical processes to also discern underlying structural equivalences or differences.

Easy: In the easier version of this task, the second graph in each pair is selected randomly from the set of all possible graphs with 8-9 nodes. Given the vast diversity within this family, non-isomorphic pairs often exhibit clear, easily discernible differences, such as varying node degrees, simplifying the detection of non-isomorphism.



Figure 4: Two non-isomorphic *degree-equivalent* graphs with the degree sequence of {3, 3, 2, 2, 1, 1}.

Hard: The hard category intensifies the challenge by selecting graph pairs that are degree-equivalent. These pairs are visually different but share identical numbers of edges and identical node degrees. Figure 4 shows an example of a non-isomorphic degree-equivalent pair of graphs.

These varied levels of difficulty are designed to progressively train and evaluate the AI's depth of understanding and its ability to perform complex visual reasoning, mirroring the cognitive processes humans use in interpreting graphical data.

3.2. Graph-Path-Finding Tasks

The Graph-Path-Finding datasets explore the model's ability to understand the concept of a path within a graph, where each graph contains 8-9 nodes. The task is divided into two subtasks:



Figure 5: Example of a Hamiltonian path, a path that visits each node a single time.

HaminItonian Path: The objective is to determine whether a graph contains a Hamiltonian path—a path that visits each vertex exactly once. Figure 5 shows an example of such a path in a graph. This is essentially a binary classification problem, with labels indicating the presence or absence of a Hamiltonian path. To maintain balance, the dataset includes an equal number of graphs with and without such paths. The training-set graphs are drawn in Kamada-Kawai layout, and the test-set graphs are drawn in planar layout. This variation in the training and testing layouts helps to assess the conceptualization ability of the models.

Shortest Path: This task presents graphs in which two

nodes are specifically marked. The goal is to find the length of the shortest path between these nodes, with potential answers of 1, 2, 3, or 4. This setup frames the task as a classification problem with four classes. The training dataset contains graphs drawn from the random layout. The test set samples are drawn from random, Kamada-Kawai, and planar layouts, forming three testing variations.

3.3. Graph-Cycle-Finding Tasks

These tasks and associated datasets test the ability of a model to understand the notion of cycles within graphs, where each graph consists of 8-9 nodes. As before, these tasks also feature two levels of difficulty, to assess different aspects of the models' capabilities.

Hamiltonian Cycle: The goal here is to determine whether the graph contains a Hamiltonian cycle, that is, a cycle which visits each vertex exactly once. The dataset is balanced between graphs that do and do-not contain a Hamiltonian cycle. To evaluate the models' ability to generalize their learning, similar to the Hamiltonian path task, the training set graphs are visualized in the Kamada-Kawai layout, whereas the test set is drawn in a planar layout.



Figure 6: An example for the biggest-chordless-cycle problem in a graph with 10 nodes.

Biggest Chordless Cycle: In this challenging scenario, the task involves identifying the size of the largest chordless cycle in a graph. A chordless cycle is defined as one where no two vertices of the cycle are connected by an edge that is not part of the cycle itself. Figure 19 shows an example of the biggest chordless cycle in a graph with 10 nodes. The possible answers are 3, 4, 5, or 6, effectively making this a classification task with four options. Similar to the previous task, the training and testing samples are drawn in the Kamada-Kawai and planar layouts, respectively.

The detailed process of the creation of the Visual-Graph-Arena datasets and example samples of each of the tasks they include is presented in the Appendix.

4. Experiments

In order to rigorously evaluate the effectiveness of Visual Graph Arena, we conducted a series of experiments, aimed at benchmarking both human and AI performance on the graph-based tasks. Initially, we established a baseline by testing human subjects, providing a comparative standard

	Isomorphism		Path	I	Cycle		
	Easy	Hard	Hamiltonian	Shortest	Hamiltonian	Biggest Chordless	
# Training Samples	140,000	127,374	25,000	80,000	69,935	80,000	
# Testing Samples	15,718	14,298	2,480	8,672	7,740	6,484	

Table 1: Number of samples in the training and test sets of the Visual Graph Arena tasks.

Table 2: Comparative performance of humans, vision models, and MLLMs on Visual Graph Arena tasks. Human participants achieved near-perfect accuracy (88.2-100%), excelling in shortest path identification. Vision models failed isomorphism tasks with ConvNeXt outperforming transformer-based models in non-failing tasks. MLLMs showed near-random performance except GPT-01, which achieved limited success in shortest path (55%) and Hamiltonian cycle (66.6%) tasks.

	Isomorphism			Path	Cycle			
	easy	hard	hamiltonian ¹	shortest path ²			hamiltonian ¹	chordless ¹
				random	kawai	planar	-	
human	95.0	91.6	96.6	98.3	100	100	93.3	88.2
Vit Base	FAIL ³	FAIL	57	67.4	66.7	64.4	67.8	32.3
Swin-T Base	FAIL	FAIL	65.8	68.6	65.7	65.4	71.1	34.4
ConvNext Base	FAIL	FAIL	72.9	73.3	82.4	73.3	74.9	36.3
SigLIP Base	54.4	FAIL	59.5	25.2	-	-	63.5	28.0
DINov2 Base	FAIL	FAIL	56.8	36.4	-	-	FAIL	31.1
GPT-o1	FAIL	FAIL	FAIL	55.0	-	-	66.0	FAIL
GPT-40	FAIL	FAIL	FAIL	FAIL	-	-	FAIL	FAIL
Claude 3.5 Sonnet	FAIL	FAIL	FAIL	FAIL	-	-	FAIL	FAIL

¹ These tasks are trained on *kawai* layout and tested on *planar* layout.

² Shortest path task is trained on *random* layout, and tested on *random*, *kawai* and *planar* layouts.

³ We use FAIL to indicate near-random performance, borrowing the term from (Tay et al., 2021).

for assessing AI capabilities in similar conditions. Following this, we moved on to evaluate advanced multimodal-llms including GPT-o1 and Claude 3.5 Sonnet, to test their proficiency in visual reasoning tasks that demand an integration of visual and textual analysis. Subsequently, we focused on the performance of leading vision models, which were specifically trained on our datasets and then tested on the designated test sets containing a different graph layout.

We detail each experiment, and discuss the methodologies employed, the models tested, and the results obtained.

Humans: To establish a baseline of human performance on the VGA tasks, we conducted experiments with 15 subjects with a total of 24 questions, distributing four questions per task. Participants were drawn from the student and staff population of the engineering faculty, possessing a general understanding of graph concepts, though not specifically in graph isomorphism, Hamiltonian paths and cycles, or chordless cycles—most requested definition and clarity of these concepts. More detail is provided in the Appendix. Vision Models: We tested three categories of vision models:

- Supervised Pre-trained Models: ViT, ConvNeXt, and Swin-T Base, each pre-trained on the ImageNet dataset using supervised learning with labeled data.
- *Multimodal Vision-Language Model*: SigLIP Base, pretrained on image-text pairs.
- *Self-Supervised Vision Model*: DINOv2 Base, trained on unlabeled image data.

All models were fine-tuned on our tasks across 10 epochs. Details regarding the training parameters, settings, and additional considerations are provided in the Appendix.

MLLMs: We evaluated GPT-40 and 01, Claude 3 Opus and 3.5 Sonnet, and Google Gemini on 100 random samples per task to determine their ability to handle graph-based visual reasoning. Unfortunately, all models performed poorly across tasks, with 01 being the only exception, showing higher than random performance in shortest path and Hamiltonian cycle problems.

4.1. Results

Table 3: Confusion matrix for GPT-o1 performance on the shortest path task. Rows represent ground-truth shortest path lengths (1–4), columns show model predictions. Values indicate prediction percentages (diagonal = correct classifications). The model performs best on length 2 (69% accuracy) but frequently confuses lengths $1\rightarrow 2$ (67%), suggesting difficulty distinguishing adjacent node path lengths.)

Output Label	1	2	3	4
Length 1	28.0	67.0	5.0	0.0
Length 2	18.0	69.0	13.0	0.0
Length 3	0.0	40.0	60.0	0.0
Length 4	0.0	5.0	40.0	55.0

Table 2 presents the results of the experiments conducted on human subjects, vision models, and MLLMs for the Visaul Graph Arena tasks. The human participants demonstrated strong performance across all tasks, achieving 100% accuracy in finding the shortest path length and over 90% accuracy in the other tasks. Finding the length of the biggest chordless cycle proved to be the most challenging for human subjects, with an accuracy of 88.2%, which was their lowest performance among all tasks.

For the vision models, we report the best validation accuracy achieved during the training epochs. Notably, all models except SigLIP failed in the isomorphism tasks, with SigLIP achieving 54.4% accuracy on the easy isomorphism task while still failing on the hard variant. This indicates a significant limitation in most models' ability to reason about graph isomorphism from visual inputs. In the other tasks, the models exhibited relatively poor performance, considering that the labels were either binary or limited to four choices. The chordless cycle task emerged as the most challenging for all vision models, with the ConvNeXt model achieving the highest accuracy of 36.3%. An interesting observation is that the ConvNeXt model consistently outperformed the Vision Transformer (ViT) and Swin Transformer (Swin-T) models on all tasks where the models did not completely fail. This performance gap was most significant in the task of finding the shortest path length in the Kawai layout, where ConvNeXt achieved an accuracy of 82.4%, surpassing ViT and Swin-T by 17.7 and 16.7 percentage points, respectively. The newer models, SigLIP and DINov2, showed mixed results, with SigLIP demonstrating some capability in isomorphism detection but both models generally underperforming compared to ConvNeXt. This finding suggests that convolutional architectures may be more effective than transformer-based models in capturing and reasoning about visual concepts in graph representations.



Figure 7: Breakdown of the performance comparison of the VIT, Swin, and ConvNeXt models to humans on the Visual Graph Arena tasks. The central area displays the performance of a random agent.

Figure 7 illustrates the performance comparison of the trained models on the tasks of the VGA benchmark. The performance of a random agent is in gray. On the graph isomorphism tasks all models performed like a random agent.

Furthermore, we evaluated the performance of advanced MLLMs on the Visual Graph Arena tasks. Quite surprisingly, most tested MLLMs, including GPT-40, Claude 3.5 Sonnet, 3 Opus, and Google Gemini, failed to solve any of the given tasks. However, GPT-o1 demonstrated partial success, achieving non-random accuracy on two specific tasks: shortest path (55% accuracy) and Hamiltonian cycle (67% accuracy). Our analysis reveals that GPT-o1's above-random performance on the Hamiltonian cycle task is primarily attributed to the model's ability to identify graphs with leaf nodes, which trivially cannot contain Hamiltonian cycles. When cases involving leaf nodes are excluded, GPTol's performance drops to near-random levels, which may also explain its failure on the related Hamiltonian path task. This suggests that the model relies on specific structural cues rather than a deeper understanding of graph theory.

For the shortest path task, GPT-o1 exhibited varying performance across different path lengths, as shown in Table 3. The confusion matrix highlights that the model performs best when the shortest path length is 2 (69% accuracy) but frequently misclassifies adjacent nodes, such as confusing lengths 1 and 2. These observations underscore the limitations of current MLLMs in handling basic graph-reasoning tasks, particularly when visual reasoning on graph-structured data is required.

The results of our experiments reveal significant challenges and limitations in the current state of the art vision models and MLLMs when it comes to reasoning about graphstructured visual data. While human participants exhibit strong performance on the Visual Graph Arena benchmark tasks, the AI models struggle to match human-level reasoning capabilities. These findings emphasize the need for further research and development of AI systems that can effectively understand and reason about abstract graphical representations, bridging the gap between human-like visual understanding and current AI capabilities. Our Visual Graph Arena may help to overcome these limitations.

5. Conceptualization Analysis in GPT-01

In our evaluation, we observe behavioral patterns in MLLMs, particularly o1, that fundamentally diverge from mature human intelligence, that questions the capability of these models to conceptualize. Specifically, we identified two distinct anomalies.

5.1. The Middle-Score Anomaly

The Middle-Score Anomaly emerges in atomic-task assessments of intelligent agents where only extreme performances (perfect or minimal) are expected, while intermediate performance is anomalous.

Consider identifying numbers ending in zero—an atomic task. For a human, we expect performance to be near either 100% or near 50% (indicating guessing, not having task knowledge). An 80% score would be anomalous because this task cannot be partially known—one either knows to check the last digit or doesn't.

5.2. The Easier-Worse Anomaly

The Easier-Worse Anomaly emerges in atomic-task assessments of intelligent agents where performance on simpler instances is unexpectedly inferior to performance on more complex variants of the same task. This pattern contradicts fundamental assumptions about cognitive capabilities.

Consider finding the shortest path between two nodes in a graph. For genuine intelligence, performing considerably better at finding shortest paths with 3 or 4 length than finding shortest path between adjacent nodes would be anomalous. Such inverted performance within a single atomic task cannot be reconciled with genuine understanding.

5.3. From Anomalies to Pseudo-Intelligence

Our experiments with GPT-o1 strongly demonstrate both anomalies. In Table 3, when tasked with finding shortest paths in graphs, the model achieves 29% accuracy for adjacent nodes (path length 1) compared to 70% accuracy for paths of length 2, a clear case of the Easier-Worse Anomaly.

The presence of both anomalies defies explanation in terms of low intelligence, but rather suggests the existence of a pseudo-intelligence (likely result of imitative-probabilistic reasoning)—a form of capability that mimics intelligence in its above-chance success rate while failing to demonstrate the conceptualization of genuine intelligence.

This limitation becomes particularly significant given the scale of contemporary training. As noted by Sutskever, current models have been trained on nearly the entire internet (Sutskever, 2024), and our investigations confirm substantial exposure to graph-related content. This brings into question whether current probabilistic training approaches are inherently incapable of achieving conceptualization. Being said, other training methods like reinforcement learning, due to their foundations in biological learning processes, may serve as a path forward for solving this challenge.

6. Conclusions

In this paper, we introduced the Visual Graph Arena, a collection of six datasets designed designed to evaluate AI models' capacity for visual reasoning and conceptual understanding through graph-based tasks. By focusing on isomorphism detection, path finding, and cycle analysis across diverse visual layouts, our work isolates the challenge of conceptualization—the ability to recognize invariant conceptual properties despite variations in representation.

Our experiments revealed performance gap between humans and state-of-the-art AI models. Humans achieved near-perfect accuracy (88–100%) across tasks, while vision models struggled significantly, especially when presented with graphs drawn using different layouts. MLLMs (GPTo1, Claude 3.5 Sonnet) struggled to perform even basic graph reasoning tasks often performing no better than random chance. This finding highlights the need for further research of AI systems that can effectively conceptualize and reason about abstract graphical representations.

The Visual Graph Arena serves as a valuable resource for the AI research community, providing a foundation for advancing the field of visual reasoning on graph-structured data. By addressing the challenges posed by this dataset, researchers can work towards developing more robust and flexible AI systems that can bridge the gap between humanlike conceptual understanding and reasoning and machine interpretation of abstract graphical information.

Future Work. Expanding visual conceptualization benchmarks to domains like chemical structures and logic circuits, where reasoning over structural invariants is critical, represents a vital next step. While developing such datasets demands domain expertise, precise annotation, and rigorous validation to ensure fidelity, these efforts will deepen our understanding of AI's conceptualization capabilities. Such datasets will accelerate progress toward AI systems capable of human-like generalization across representations.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

Acknowledgments

Zahra Babaiee and Radu Grosu are supported by the Austrian Science Fund (FWF) project MATTO-GBM I 6605. Peyman M. Kiasari is supported by the TU Wien TrustACPS PhD School program that is supported by TTTech Auto and B&C Privatstiftung.

References

- Abboud, R., Ceylan, I. I., Grohe, M., and Lukasiewicz, T. The surprising power of graph neural networks with random node initialization. *arXiv preprint arXiv:2010.01179*, 2020.
- Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Zitnick, C. L., and Parikh, D. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pp. 2425–2433, 2015.
- Balcilar, M., Héroux, P., Gauzere, B., Vasseur, P., Adam, S., and Honeine, P. Breaking the limits of message passing graph neural networks. In *International Conference on Machine Learning*, pp. 599–608. PMLR, 2021.
- Chen, Z., Chen, L., Villar, S., and Bruna, J. Can graph neural networks count substructures? *Advances in neural information processing systems*, 33:10383–10395, 2020.
- Cherian, A., Peng, K.-C., Lohit, S., Matthiesen, J., Smith, K. A., and Tenenbaum, J. B. Evaluating large vision-andlanguage models on children's mathematical olympiads. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024. URL https://openreview.net/forum? id=YMAU2kJgzY.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houlsby, N. An image is worth 16x16 words: Transformers for image recognition at scale, 2020.
- Girshick, R., Donahue, J., Darrell, T., and Malik, J. Regionbased convolutional networks for accurate object detection and segmentation. *IEEE transactions on pattern* analysis and machine intelligence, 38(1):142–158, 2015.
- Gong, B., Tan, S., Feng, Y., Xie, X., Li, Y., Chen, C., Zheng, K., Shen, Y., and Zhao, D. UKnow: A unified

knowledge protocol with multimodal knowledge graph datasets for reasoning and vision-language pre-training. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024. URL https://openreview.net/forum? id=LdRZ9SFBku.

- Goyal, Y., Khot, T., Summers-Stay, D., Batra, D., and Parikh, D. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 6904–6913, 2017.
- Guetta, N. B., Slobodkin, A., Maimon, A., Habba, E., Rassin, R., Bitton, Y., Szpektor, I., Globerson, A., and Elovici, Y. Visual riddles: a commonsense and world knowledge challenge for large vision and language models. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024. URL https://openreview.net/forum? id=hwbRjslR5N.
- Gupta, T. and Kembhavi, A. Visual programming: Compositional visual reasoning without training. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 14953–14962, 2023.
- Han, C., Mao, J., Gan, C., Tenenbaum, J., and Wu, J. Visual concept-metaconcept learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- Haresh, S., Dijkman, D., Bhattacharyya, A., and Memisevic, R. Clevrskills: Compositional language and visual reasoning in robotics. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024. URL https://openreview. net/forum?id=64sZtFSOh6.
- Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., Catasta, M., and Leskovec, J. Open graph benchmark: Datasets for machine learning on graphs. *Advances in neural information processing systems*, 33:22118–22133, 2020.
- Johnson, J., Hariharan, B., Van Der Maaten, L., Fei-Fei, L., Lawrence Zitnick, C., and Girshick, R. 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.
- Kahou, S. E., Michalski, V., Atkinson, A., Kádár, Á., Trischler, A., and Bengio, Y. Figureqa: An annotated figure dataset for visual reasoning. *arXiv preprint arXiv:1710.07300*, 2017.

- Kamada, T., Kawai, S., et al. An algorithm for drawing general undirected graphs. *Information processing letters*, 31(1):7–15, 1989.
- Li, Y., Hu, B., Shi, H., Wang, W., Wang, L., and Zhang, M. Visiongraph: Leveraging large multimodal models for graph theory problems in visual context. *arXiv preprint arXiv:2405.04950*, 2024.
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., and Guo, B. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (*ICCV*), pp. 10012–10022, October 2021.
- Liu, Z., Mao, H., Wu, C.-Y., Feichtenhofer, C., Darrell, T., and Xie, S. A convnet for the 2020s. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- Ma, Y., Zang, Y., Chen, L., Chen, M., Jiao, Y., Li, X., Lu, X., Liu, Z., Ma, Y., Dong, X., Zhang, P., Pan, L., Jiang, Y.-G., Wang, J., Cao, Y., and Sun, A. MMLONGBENCH-DOC: Benchmarking long-context document understanding with visualizations. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024. URL https://openreview.net/forum? id=loJMlacwzf.
- Morris, C., Kriege, N. M., Bause, F., Kersting, K., Mutzel, P., and Neumann, M. Tudataset: A collection of benchmark datasets for learning with graphs. *arXiv preprint arXiv:2007.08663*, 2020.
- Murphy, R., Srinivasan, B., Rao, V., and Ribeiro, B. Relational pooling for graph representations. In *International Conference on Machine Learning*, pp. 4663–4673. PMLR, 2019.
- Shao, H., Qian, S., Xiao, H., Song, G., Zong, Z., Wang, L., Liu, Y., and Li, H. Visual cot: Advancing multi-modal language models with a comprehensive dataset and benchmark for chain-of-thought reasoning. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024. URL https: //openreview.net/forum?id=aXeiCbMFFJ.
- Sutskever, I. Sequence to sequence learning with neural networks: What a decade. Test of Time Award talk at the 38th Conference on Neural Information Processing Systems (NeurIPS 2024), 2024. Available at https: //www.youtube.com/watch?v=HlGi400uZyw.
- Tay, Y., Dehghani, M., Abnar, S., Shen, Y., Bahri, D., Pham, P., Rao, J., Yang, L., Ruder, S., and Metzler, D. Long range arena: A benchmark for efficient transformers. In

International Conference on Learning Representations, 2021. URL https://openreview.net/forum? id=qVyeW-grC2k.

- Wang, K., Pan, J., Shi, W., Lu, Z., Ren, H., Zhou, A., Zhan, M., and Li, H. Measuring multimodal mathematical reasoning with MATH-vision dataset. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024. URL https: //openreview.net/forum?id=QWTCcxMpPA.
- Wang, Y. and Zhang, M. Towards better evaluation of gnn expressiveness with brec dataset. arXiv preprint arXiv:2304.07702, 2023.
- Yang, Z., Gan, Z., Wang, J., Hu, X., Lu, Y., Liu, Z., and Wang, L. An empirical study of gpt-3 for few-shot knowledge-based vqa. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 3081–3089, 2022.

A. Dataset Description

The Visual Graph Arena (VGA) Benchmark is a collection of six datasets designed to evaluate and enhance the visual reasoning capabilities of AI models on graph-based tasks. The benchmark aims to bridge the gap between human-like understanding and machine interpretation of abstract graphical information. The datasets are intended for research purposes in the field of artificial intelligence, specifically focusing on advancing visual reasoning on graph-structured data.

The benchmark datasets are hosted on the Visual Graph Arena website at visualgrapharena.com. The website provides access to the datasets, along with documentation and usage guidelines.

Figure 8 illustrates the directory structure of the Visual Graph Arena (VGA) dataset. The VGA dataset is structured to facilitate easy access to various graph analysis tasks including isomorphism, path problems, and cycle detection. Each category is further subdivided into training and testing datasets with detailed sample counts and label files. For each input data, 'image_{i}.png' corresponds to the i-th row of it's 'label.txt'. For more information and to access the dataset, please visit the VGA dataset website at visualgrapharena.com.

Author Statement: We, the authors, take full responsibility for any potential violation of rights related to the Visual Graph Arena dataset and its materials. We confirm that the dataset is our original work and grant the hosting platform permission to distribute it under the CC BY 4.0 license. We will ensure proper access and maintenance of the dataset.

B. Dataset Creation

To create the datasets, we utilized the collection of connected graphs with 8 and 9 nodes from https://users.cecs. anu.edu.au/~bdm/data/graphs.html. This collection comprises 11,117 connected graphs with 8 nodes and 261,080 connected graphs with 9 nodes. The Python library NetworkX was employed to visualize the graphs in plots for the datasets. Each sample was plotted within a square canvas of 700×700 pixels. The following details the specific steps involved in creating each task:

B.1. Isomorphism Tasks

Easy Isomorphism: For the easy isomorphism tasks, we plotted two graphs in a single plot, with the graph on the right using the Kamada-Kawai layout and the graph on the left using the planar layout. As not all graphs are planar, we first identified and selected all planar graphs from the set of connected graphs. We then split these planar graphs into separate test and train sets, ensuring that no graph from the training set appeared in the test set. The true and false samples were created as follows:

- True Labels: For each true label, we took a graph and plotted it using both the Kamada-Kawai and planar layouts.
- False Labels: For each false label, we plotted one graph using the Kamada-Kawai layout and then plotted another randomly selected graph from the set using the planar layout.

To maintain a balanced dataset, we ensured an equal number of true and false labels in both the training and test sets. Figure 9 presents examples of the easy isomorphism task with both true and false labels.

Hard Isomorphism: For this task, we proceeded similarly to the previous one, but with a key difference in the selection of graphs for false labels. Instead of picking a random graph from the set, we searched for graphs within the set that are degree-equivalent to the first graph. From these degree-equivalent graphs, we randomly chose one and plotted it using the planar layout. In cases where no degree-equivalent graphs were found for the first graph within the set, we discarded that sample. As a result, the number of samples in this task is smaller compared to the easy isomorphism task. Figure 10 presents examples of the hard isomorphism task with both true and false labels.

B.2. Path Finding Tasks

Shortest Path: For the shortest path task, we created a balanced set of graphs with shortest paths of lengths 1, 2, 3, and 4 between two marked nodes. We opted not to include graphs with shortest paths of length 5 or longer, as the number of such graphs within the 8-9 node graph set was too small to maintain a balanced dataset with an equal number of samples for each path length. We began by separating the planar graphs for the test set. From the remaining non-planar graphs, which were

used for the training set, we identified pairs of nodes with shortest paths of lengths 1, 2, 3, and 4. These nodes were marked with squares in the plot, while the rest of the nodes were represented as circles. The training set graphs were plotted using the Kamada-Kawai layout. For each path length, we plotted 20,000 samples, resulting in a total of 80,000 samples for the training set. For the test set, we created a similarly balanced set of graphs with different path lengths and plotted them using the planar layout. Figure 11 shows examples of the shortest path task.

Hamiltonian Paths: For this task, we first identified the graphs that contain a Hamiltonian path from the entire set of 8-9 connected graphs. To create a balanced dataset with an equal number of true and false labels, we then randomly selected the same number of graphs from the remaining non-Hamiltonian graphs as the number of Hamiltonian graphs found. This step was necessary since Hamiltonian graphs are fewer in count compared to non-Hamiltonian graphs. After creating the balanced dataset, we split the set into train and test subsets. For the training set, we plotted the graphs using the Kamada-Kawai layout, while for the test set, we used the planar layout. Figure 12 shows examples of the Hamiltonian path task.

B.3. Cycle Finding Tasks

Hamiltonian Cycles: The task of finding Hamiltonian cycles was approached in a similar manner to the Hamiltonian path task. We first identified graphs that contain a Hamiltonian cycle from the entire set of 8-9 connected graphs. To create a balanced dataset with an equal number of true and false labels, we randomly selected the same number of graphs from the remaining non-Hamiltonian cycle graphs as the number of Hamiltonian cycle graphs found. After creating the balanced dataset, we split the set into train and test subsets. For the training set, we plotted the graphs using the Kamada-Kawai layout, while for the test set, we used the planar layout. Figure 13 presents examples of the Hamiltonian cycle task.

Chordless Cycles: The chordless cycle task was similar to the shortest path task, but instead focused on the largest chordless cycle lengths of 3, 4, 5, and 6. We chose not to include graphs with no chordless cycles (trees) because the number of such graphs was too small to maintain a balanced set.

We first separated the planar graphs for the test set. From the remaining non-planar graphs, which were used for the training set, we identified the largest chordless cycles of lengths 3, 4, 5, and 6. For each chordless cycle length, we plotted an equal number of samples, resulting in a balanced dataset. The training set graphs were plotted using the Kamada-Kawai layout.

For the test set, we created a similarly balanced set of graphs with different chordless cycle lengths and plotted them using the planar layout. Figure 19 shows examples of the chordless cycles task.

C. Training the Vision Models

For training the vision models on the Visual Graph Arena benchmark tasks, we utilized three state-of-the-art architectures: ConvNeXt base, Swin Transformer base, and ViT (Vision Transformer) base. All models were initially pre-trained on the ImageNet dataset and then fine-tuned on our specific graph-based tasks. The training process was implemented using PyTorch, with the Adam optimizer and a learning rate of 1e - 4. The batch size was set to 32, and the models were trained for 10 epochs on each task. The CrossEntropyLoss was used as the loss function, as the tasks were formulated as either binary or multi-class classification problems. We used 2 NVIDIA TITAN RTX GPUs for training the models.

The input images were preprocessed using transformations from the PyTorch transforms module. The images were resized to 384×384 pixels, and the pre-trained models were trained on the same size of images on ImageNet.

To handle the different output requirements of the tasks, the number of output classes in the final layer of the models was adjusted accordingly. For the isomorphism tasks and the Hamiltonian path/cycle tasks, the number of output classes was set to 2 (binary classification). For the shortest path and chordless cycle tasks, the number of output classes was set to 4, corresponding to the different path lengths or cycle sizes.

During training, the model's performance was evaluated on the test set at the end of each epoch. The test set consisted of graphs drawn using different layouts than those used in the training set, allowing us to assess the model's ability to generalize to different visual representations. After the training process, the best-performing model on the test set was selected for each task and architecture combination. The code for training the vision models, written in Python using PyTorch, is provided in the supplementary materials and will be publicly available in a GitHub repository.



Figure 8: Overview of the directory structure of the Visual Graph Arena benchmark.



Figure 9: Examples of the easy isomorphism task with true and false labels.



Figure 10: Examples of the easy isomorphism task with true and false labels.



Figure 11: Examples of the shortest path task between marked nodes.



Figure 12: Examples of the Hamiltonian path task with true and false labels.



Figure 13: Examples of the Hamiltonian cycle task with true and false labels.



(a) The biggest chordless cycle is of length 6.

(b) The biggest chordless cycle is of length 4.



D. MLLMs adjacency matrix creation

In this section, we evaluate the performance of GPT-40 and 01, Claude 3 Opus, and Claude 3.5 Sonnet in transcribing graph structures. We presented these models with three graph images and provided the prompt: "The image presented contains a graph. Please transcribe this graph's adjacency matrix." We then analyzed their ability to generate accurate adjacency matrices from the visual input. Our results reveal that all models struggle to correctly transcribe the graphs, with varying degrees of accuracy in extracting basic graph properties. GPT-4 comes close to accurately identifying the number of nodes but struggles with edge counts. Claude 3.5 Sonnet shows the best performance, closely approximating both the number of nodes and edges. In contrast, Claude 3 Opus performs poorly across all metrics. We visualized the graphs transcribed by these models to assess their accuracy. Although Claude 3.5 Sonnet was not available at the time of our initial submission, we have incorporated its responses into our analysis, because it demonstrates superior performance compared to both GPT-4 and Claude 3 Opus, showing a marked improvement over Opus in particular.



Figure 15: The three graphs we asked the LLMs to generate the adjacency matrix.



Figure 16: reconstructed graphs from the generated adjacency matrix by GPT-4o.

E. LLMs Prompts and Responses Samples

In this section, we present one example of each task, along with the prompts used and the responses from GPT-01, GPT-4, Claude 3 Opus, and Claude 3.5 Sonnet. We highlight correct final answers in green and incorrect ones in red. The results demonstrate that these models can analyze graphs, provide chains of thought, and attempt to track nodes and edges. However,



(a) Graph a

Figure 17: reconstructed graphs from the generated adjacency matrix by Claude 3.5 Sonnet.



Figure 18: reconstructed graphs from the generated adjacency matrix by Claude 3 Opus.

they often fail to solve the problems accurately. In the example in three cases, while the final yes/no responses are correct, closer inspection of the analysis reveals obvious mistakes, such as incorrectly counting the number of nodes. The correct responses are not surprising given the yes/no nature of some tasks. Notably, in the shortest-path task, all models unexpectedly fail to detect that the two marked nodes are connected by a single edge.

E.1. Isomorphism - Easy

prompt: The image presented contains two side-by-side graphs. Without using code, please answer whether the two given graphs are isomorphic. Please report your final answer as 'Yes' or 'No' as your very last word.



Figure 19: reconstructed graphs from the generated adjacency matrix by GPT o1.

GPT-40 Output:

To determine if the two graphs are isomorphic, we need to check whether there is a one-to-one correspondence between the vertices of the two graphs that preserves the adjacency relationships. Looking at the two graphs in the image:

- 1. Number of vertices: Both graphs have the same number of vertices.
- 2. Number of edges: Both graphs also have the same number of edges.
- 3. Vertex degrees: The degrees of corresponding vertices (number of edges connected to each vertex) should be the same.

However, upon visual inspection, the structure of the graphs seems to differ significantly. The graph on the right has a more symmetrical structure, while the graph on the left appears more irregular. Specifically, the arrangement of connections and the overall shape of the graphs suggest that their connectivity patterns differ. This implies that there isn't a one-to-one correspondence that preserves the adjacency relationships between the vertices of the two graphs. Therefore, the final answer is **No**.



Figure 20: The input image, the answer is 'Yes'

Claude 3.5 Sonnet Output:

To determine if these graphs are isomorphic, I'll analyze their structures: Left graph:

- 9 vertices
- 10 edges
- Degree sequence: 3, 3, 3, 2, 2, 2, 2, 2, 1

Right graph:

- 9 vertices
- 10 edges
- Degree sequence: 4, 3, 2, 2, 2, 2, 2, 1

While both graphs have the same number of vertices and edges, their degree sequences are different. The left graph has a maximum degree of 3, while the right graph has a vertex with degree 4. Isomorphic graphs must have identical degree sequences. Since these graphs have different degree sequences, they cannot be isomorphic.

No

Claude 3 Opus Output:

To determine if the two graphs are isomorphic, I'll check if there exists a bijective mapping between their vertex sets that preserves adjacency.

Both graphs have 5 vertices. Let's label the vertices of the left graph L1 to L5 from top to bottom, and the vertices of the right graph R1 to R5 from top to bottom.

Comparing vertex degrees:

L1, L2, L4, L5 have degree 3. L3 has degree 4.

R1, R3, R5 have degree 3. R2, R4 have degree 4.

Since the degree sequences don't match, there cannot be an adjacency-preserving bijection between the vertex sets. Therefore, the two graphs are not isomorphic.

No.

GPT-01 Output:

Thought about isomorphic graphs for 3m 21s

A quick way to see they cannot match is by comparing vertex-degrees. In the right-hand graph, the "apex" vertex (the topmost point) clearly has degree 5 (it has edges to five other vertices). In the left-hand graph, no vertex has that many incident edges (each has degree at most 3 or 4). Since an isomorphism would require a one-to-one match of vertices with the same degree, the two graphs cannot be isomorphic.

E.2. Isomorphism - Hard

prompt: The image presented contains two side-by-side graphs. Without using code, please answer whether the two given graphs are isomorphic. Please report your final answer as 'Yes' or 'No' as your very last word.



Figure 21: The input image, the answer is 'No'

GPT-4o Output:

To determine if the two graphs are isomorphic, I will analyze the structural properties of each graph, such as the number of vertices, edges, and the degree of each vertex, and compare them. **Step 1: Count the vertices and edges.**

- Left graph:
 - Number of vertices: 10
 - Number of edges: 13
- Right graph:
 - Number of vertices: 10
 - Number of edges: 13

Step 2: Compare the degree of each vertex.

• Both graphs should have the same degree sequence if they are isomorphic. The degree of each vertex is the number of edges connected to it.

Step 3: Visual inspection.

• Compare the overall structure and connectivity to see if a one-to-one correspondence exists between the vertices and edges of the two graphs.

Upon close inspection, I can conclude that both graphs have the same number of vertices and edges, and they appear to have a similar degree sequence and connectivity structure. Given these observations, the graphs seem to be isomorphic.

Final answer: Yes .

Claude 3.5 Sonnet Output:

To determine if the two graphs are isomorphic, I'll analyze their structures: Left graph:

- 11 vertices
- 15 edges
- Complex structure with multiple triangles and intersecting lines

Right graph:

- 11 vertices
- 15 edges
- More elongated structure, fewer triangles, more "tree-like" in parts

While both graphs have the same number of vertices and edges, their structures appear quite different:

- 1. The left graph is more densely interconnected, with multiple triangles and a more compact shape.
- 2. The right graph has a more stretched out, less symmetrical structure.
- 3. The degree distribution (number of edges connected to each vertex) seems different between the two graphs.
- 4. The left graph appears to have more cycles than the right graph.

These structural differences suggest that there is no one-to-one mapping that could transform one graph into the other while preserving all connections.

Given these observations, my conclusion is that these graphs are not isomorphic.

No

Claude 3 Opus Output:

Okay, let's carefully analyze the two graphs to determine if they are isomorphic. To be isomorphic, the two graphs must have:

- 1. The same number of vertices
- 2. The same number of edges
- 3. Equivalent connectivity between corresponding pairs of vertices

Comparing the two given graphs:

- 1. Number of vertices: Both graphs appear to have 6 vertices each. So this condition is satisfied.
- 2. Number of edges: I count 9 edges in the graph on the left, and 9 edges in the graph on the right as well. Edge count matches.
- 3. Vertex connectivity: This is trickier to verify visually. We'd need to find a one-to-one mapping between the vertices of the two graphs, such that the edge connections are preserved.

One potential mapping: Left graph vertex (from top) $1 \rightarrow 1$, $2 \rightarrow 2$, $3 \rightarrow 6$, $4 \rightarrow 5$, $5 \rightarrow 4$, $6 \rightarrow 3$ on right graph This mapping appears to preserve all the edge connections between corresponding vertex pairs. Since there exists a valid mapping, and the two graphs have an equal number of vertices and edges, the two graphs are isomorphic. Therefore, my final answer is Yes.

GPT-o1 Output:

Thought about graph isomorphism for 3m 45s

A quick way to see they differ is by looking at how each drawing "uses" its outer boundary. In the right-hand graph, all vertices lie on or inside a single large triangular boundary; in the left-hand graph, the outer boundary clearly has more than three vertices. That alone forces a mismatch in adjacency that no relabeling can fix.

Hence, the two graphs are not isomorphic.

No

E.3. Cycle - Hamiltonian

prompt: The image presented contains a graph. Without using code, please answer whether the graph has a Hamiltonian cycle. Please report your final answer as 'Yes' or 'No' as your very last word.



Figure 22: The input image, the answer is 'Yes'

GPT-40 Output:

To determine whether the graph has a Hamiltonian cycle (a cycle that visits every vertex exactly once and returns to the starting vertex), we need to consider the structure of the graph and see if such a cycle can be traced. Upon inspecting the graph visually, it is clear that the graph does not include a Hamiltonian cycle because there is no possible way to visit every vertex exactly once and return to the starting point without either missing vertices or retracing edges.

The final answer is **No**.

Claude 3.5 Sonnet Output:

To determine if this graph has a Hamiltonian cycle, we'll examine if there's a path that visits each vertex exactly once and returns to the starting vertex.

The graph has 9 vertices connected by lines.

A Hamiltonian cycle would need to:

- 1. Visit all 9 vertices
- 2. Use only existing edges
- 3. Return to the starting point
- 4. Visit each vertex only once (except the start/end)

Looking at the graph structure:

- There are several vertices with only two connections, which limits path options.
- The graph is not fully connected it has a "pinched" middle section.
- There's no obvious circular path that covers all vertices.

Attempting to trace potential paths:

- Starting from any outer vertex, it's impossible to reach all other vertices and return without revisiting some.
- The central "pinched" area forces any potential cycle to cross over itself, which isn't allowed in a Hamiltonian cycle.

Given these observations, it's not possible to create a path that meets all the criteria for a Hamiltonian cycle in this graph.

No

Claude 3 Opus Output:

To determine if the given graph has a Hamiltonian cycle, we need to check if there exists a cycle that visits each vertex exactly once and returns to the starting vertex. Looking at the graph, we can observe the following:

- 1. The graph has 6 vertices.
- 2. Each vertex has a degree (number of edges connected to it) of at least 2.
- 3. It appears that we can trace a path that visits each vertex exactly once and returns to the starting point.

One possible Hamiltonian cycle in this graph could be:

Start at any vertex, move to an adjacent vertex, then continue moving to unvisited adjacent vertices until all vertices have been visited once, and finally return to the starting vertex.

Therefore, based on the observation that a cycle visiting each vertex exactly once exists in the given graph, my final answer is Yes.

GPT-o1 Output:

Thought about Hamiltonian cycle in graph for 3m 47s From a purely visual/graph-theoretic inspection (tracking each vertex and checking that we can "thread" a single cycle through all of them exactly once), one finds that there is indeed a way to visit all vertices in a cycle. In other words, the graph is Hamiltonian. Yes

E.4. Cycle - Biggest Chordless

prompt: The image presented contains a graph. Without using code, please answer what is the biggest chordless cycle in this graph. Please report your final answer as number in your very last word.



Figure 23: The input image, the answer is '3'

GPT-40 Output:

To find the biggest chordless cycle in the given graph, let's analyze the connections between the vertices step by step. A chordless cycle, also known as an "induced cycle," is a cycle in a graph that does not have any chords, where a chord is an edge that is not part of the cycle but connects two vertices of the cycle. **Steps to Analyze the Graph:**

- 1. Identify all cycles: We need to identify all the possible cycles in the graph.
- 2. **Determine if they are chordless:** For each cycle, check if there are any extra edges (chords) connecting non-consecutive vertices in the cycle.

Observation from the Image: Looking at the graph, a few potential cycles stand out:

- Cycle 1: Going through five points that seem to form a pentagon shape without any internal edges connecting non-adjacent vertices.
- Cycle 2: A cycle involving four vertices.

Upon inspecting the graph:

- 1. The cycle involving the five vertices appears to be the largest cycle.
- 2. There are no additional edges (chords) connecting non-consecutive vertices in this cycle, making it a chordless cycle.

Conclusion: The biggest chordless cycle in this graph consists of **5 vertices**. Thus, the final answer is **5**.

Claude 3.5 Sonnet Output:

To find the biggest chordless cycle in this graph, I'll examine the structure carefully:

- 1. First, I'll identify all cycles in the graph.
- 2. Then, I'll check each cycle to see if it has any chords (edges connecting non-adjacent vertices within the cycle).
- 3. Among the chordless cycles, I'll determine the largest one.

After analyzing the graph:

- There are several triangles (3-cycles), but these are too small to be the answer.
- I can see a few 4-cycles, but upon closer inspection, they all have chords.
- The largest cycle I can find without chords is a 5-cycle.

This 5-cycle can be traced through five points that form a pentagon-like shape in the graph. Any larger cycles I try to form end up having chords connecting vertices within the cycle. Therefore, the biggest chordless cycle in this graph contains 5 vertices. To express my final answer as requested: 5

Claude 3 Opus Output:

To find the largest chordless cycle in the given graph, I'll analyze the graph visually. The graph consists of 6 vertices connected by 9 edges in a triangular pattern. I can see several cycles:

- The outermost cycle connecting all 6 vertices
- Two smaller cycles, each connecting 4 vertices, sharing one edge of the outer cycle
- Four cycles connecting 3 vertices each (the smallest cycles possible)

The 3-vertex cycles all have chords (edges connecting non-adjacent vertices in the cycle), so they are not chordless. The 4-vertex cycles don't have any chords. These are the largest chordless cycles in the graph. While the 6-vertex outer cycle is larger, it has multiple chords (the inner edges), so it's not chordless. Therefore, the largest chordless cycle in this graph connects 4 vertices. 4

GPT-o1 Output:

Thought about chordless cycle size for 2m 30s

A succinct way to spot the largest chord-free ("induced") cycle is to look around the "outer rim" of the drawing and check whether any diagonals cut across those vertices. By inspection, you will see that all 6 vertices forming the perimeter loop in the lower and middle/top portions can be traversed in a cycle without encountering an internal chord. Any attempt to include the interior edges or to form a 7-cycle will run into an extra cross-edge (a chord). Hence the longest chordless cycle in this graph has length **6**.

E.5. Path - Hamiltonian

prompt: The image presented contains a graph. Without using code, please answer whether the graph has a Hamiltonian path. Please report your final answer as 'Yes' or 'No' as your very last word.



Figure 24: The input image, the answer is 'Yes'

GPT-4o Output:

To determine if the given graph has a Hamiltonian path, I will visually inspect the graph. A Hamiltonian path is a path in a graph that visits each vertex exactly once. Upon inspecting the graph:

- 1. I identified the vertices and observed the connections between them.
- 2. I checked if there's a path that visits every vertex exactly once.

Given the structure of the graph, it appears that there is **no** such path that visits every vertex exactly once without retracing any edge or missing any vertex. Final answer: **No**

Claude 3.5 Sonnet Output:

To determine if this graph has a Hamiltonian path, we'll examine whether there's a path that visits each vertex exactly once.

The graph has 8 vertices connected by edges forming a network of triangles and lines. Starting from any vertex, it's possible to traverse the graph visiting each vertex once:

- 1. Begin at the leftmost point
- 2. Move right to the next point
- 3. Go up to the top of the leftmost triangle
- 4. Move right to the center point of the graph
- 5. Go down to the bottom point of the rightmost triangle
- 6. Move up to the right point of that triangle
- 7. Go to the top point of the rightmost triangle
- 8. End at the uppermost point of the graph

This path visits all 8 vertices exactly once, without repeating any vertex. Therefore, the graph does have a Hamiltonian path. Yes

Claude 3 Opus Output:

To determine whether the given graph has a Hamiltonian path, I'll analyze its structure and properties. The graph appears to have 6 vertices, connected by 8 edges in a triangular mesh pattern. Every vertex has a degree (number of edges connected to it) of either 2 or 4.

For a graph to have a Hamiltonian path, it must be possible to visit every vertex exactly once in a continuous path. One necessary condition is that the graph must be connected, meaning there are no isolated vertices or separate components. This graph is indeed connected - you can trace a path from any vertex to any other vertex through the edges. This satisfies the connectivity requirement.

However, the graph's structure poses a challenge. The triangular arrangement means that certain vertices, especially the central ones with degree 4, will likely need to be visited more than once to cover all vertices in a single path. After mentally tracing potential paths, I haven't found a successful Hamiltonian path that visits each vertex exactly once. The graph's layout seems to necessitate revisiting some vertices. Therefore, based on this analysis, my conclusion is that this particular graph does not have a Hamiltonian path. No.

GPT-o1 Output:

Thought about Hamiltonian path in graph for 2m 49s

From inspection, one can find a route that visits each vertex exactly once (in particular, the "lone" degree-1 vertex at the top can serve as one end of such a path). Therefore the graph does indeed admit a Hamiltonian path. Yes

E.6. Path - Shortest

prompt: The image presented contains a graph. Two nodes are marked by square. Without using code, please answer what is the shortest path between the two square nodes in this graph. Please report your final answer as number in your very last word.



Figure 25: The input image, the answer is '1'

GPT-40 Output:

To find the shortest path between the two square nodes in the graph, we can visually inspect the connections between the nodes.

Looking at the image:

- 1. Start from the square node on the left side of the graph.
- 2. There is a direct connection to the node above it.
- 3. From this node, there is a connection to the square node on the right side.

Therefore, the shortest path consists of **two edges**. The final answer is **2**.

Claude 3.5 Sonnet Output:

To find the shortest path between the two square nodes in this graph, I'll count the minimum number of edges that need to be traversed to get from one square node to the other. Looking at the graph, I can see that there are two potential paths:

- 1. A direct path that goes through the center node of the graph.
- 2. A path that follows the outer edges of the graph.

The direct path through the center is clearly shorter. It requires traversing only 2 edges:

- 1. From the first square node to the center node
- 2. From the center node to the second square node

The path along the outer edges would require traversing at least 3 edges. Therefore, the shortest path between the two square nodes in this graph is 2.

Claude 3 Opus Output:

The image shows a graph with 6 nodes, where 2 nodes are marked by squares. To find the shortest path between the two square nodes:

- 1. Starting from the top square node, there are two possible paths: down-left or down-right.
- 2. Following the down-left path:

Square \rightarrow Circle \rightarrow Square

This path has a length of 2 edges.

3. Following the down-right path:

Square \rightarrow Circle \rightarrow Circle \rightarrow Square

This path has a length of 4 edges.

Therefore, the shortest path between the two square nodes goes through the down-left path, with a length of 2 edges. The shortest path between the two square nodes is 2.

GPT-o1 Output:

Thought about shortest path in graph for 3m 21s

From inspection of the diagram's edges, there is no direct edge connecting the two square-marked nodes; however, they lie exactly two edges apart (via a single intermediate vertex).

2