

# Novelty Assessment Report

**Paper:** Chart Deep Research in LVLMs via Parallel Relative Policy Optimization

**PDF URL:** <https://openreview.net/pdf?id=cRJNk4bhZi>

**Venue:** ICLR 2026 Conference Submission

**Year:** 2026

**Report Generated:** 2026-01-01

## Abstract

With the rapid advancement of data science, charts have evolved from simple numerical presentation tools to essential instruments for insight discovery and decision-making support. However, current chart data intelligence exhibits significant limitations in deep research capabilities, with existing methods predominantly addressing shallow tasks such as visual recognition or factual question-answering, rather than the complex reasoning and high-level data analysis that deep research requires. This limitation stems from two primary technical bottlenecks: at the training level, existing post-training techniques exhibit deficiencies in handling multi-dimensional reward signal interference and heterogeneous data gradient conflicts, preventing models from achieving balanced development across multiple capability dimensions; at the evaluation level, current methods remain limited to factual retrieval and basic computation, failing to assess end-to-end analytic reasoning and other deep research capabilities. To address the training challenge, we propose PRPO, which performs parallel optimization across reward dimensions and capability partitioning across data types, effectively disentangling conflicts between heterogeneous data and multi-dimensional reward signals while ensuring optimization stability. For the evaluation challenge, we construct MCDR-Bench based on the "error uniqueness principle," transforming subjective generation assessment into objective error identification through controllable error injection, enabling quantifiable evaluation of deep research capabilities. Experimental validation confirms that the proposed PRPO and MCDR-Bench jointly establish a unified framework that systematically advances chart deep research through enhanced collaborative training and objective evaluation.

### Disclaimer

This report is **AI-GENERATED** using Large Language Models and WisPaper (a scholar search engine). It analyzes academic papers' tasks and contributions against retrieved prior work. While this system identifies **POTENTIAL** overlaps and novel directions, **ITS COVERAGE IS NOT EXHAUSTIVE AND JUDGMENTS ARE APPROXIMATE**. These results are intended to assist human reviewers and **SHOULD NOT** be relied upon as a definitive verdict on novelty.

Note that some papers exist in multiple, slightly different versions (e.g., with different titles or URLs). The system may retrieve several versions of the same underlying work. The current automated pipeline does not reliably align or distinguish these cases, so human reviewers will need to disambiguate them manually.

If you have any questions, please contact: mingzhang23@m.fudan.edu.cn

## Core Task Landscape

This paper addresses: **Chart Analysis Using Large Vision-Language Models**

A total of **50 papers** were analyzed and organized into a taxonomy with **23 categories**.

### Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Model Architecture and Training Methodologies**
- **Reasoning and Inference Mechanisms**
- **Evaluation and Benchmarking**
- **Application Domains and Task-Specific Adaptations**
- **Foundational Studies and Surveys**

### Complete Taxonomy Tree

- Chart Analysis Using Large Vision-Language Models Survey Taxonomy
- Model Architecture and Training Methodologies
  - Vision Vocabulary and Encoding Enhancement (1 papers)
  - [7] Vary: Scaling up the Vision Vocabulary for Large Vision-Language Models (Wei, 2023) [View paper](#)
  - Pre-training Objectives and Cross-modal Alignment (2 papers)
  - [49] In-Depth and In-Breadth: Pre-training Multimodal Language Models Customized for Comprehensive Chart Understanding (Fan, 2025) [View paper](#)
  - [50] Enhanced Chart Understanding in Vision and Language Task via Cross-modal Pre-training on Plot Table Pairs (ZHOU Mingyang, 2023) [View paper](#)
  - Post-training Optimization and Reinforcement Learning ★ (1 papers)
  - [0] Chart Deep Research in LVLMs via Parallel Relative Policy Optimization (Anon et al., 2026) [View paper](#)
  - Efficient Model Design and Compression (2 papers)
  - [8] Tinychart: Efficient chart understanding with visual token merging and program-of-thoughts learning (Zhang Liang, 2024) [View paper](#)
- Reasoning and Inference Mechanisms
  - Chain-of-Thought and Step-by-Step Reasoning (1 papers)
  - [2] End-to-end chart summarization via visual chain-of-thought in vision-language models (Raymond Choi, 2025) [View paper](#)
  - Code-Driven Reasoning and Program Synthesis (2 papers)
  - [47] ChartGen: Scaling Chart Understanding Via Code-Guided Synthetic Chart Generation (Kondic, 2025) [View paper](#)
  - [48] Chart-CoCa: Self-Improving Chart Understanding of Vision LMs via Code-Driven Synthesis and Candidate-Conditioned Answering (Gongyao Jiang, 2025) [View paper](#)
  - Multimodal Feedback and Iterative Refinement (3 papers)
  - [4] Closing the feedback loop in text2vis: Refining visualization with vision-language models (Shengze Shi, 2025) [View paper](#)
  - [24] ChartSketcher: Reasoning with Multimodal Feedback and Reflection for Chart Understanding (Zhang Ling-ling, 2025) [View paper](#)
  - [33] Socratic chart: Cooperating multiple agents for robust svg chart understanding (Ji, 2025) [View paper](#)

- Evaluation and Benchmarking
  - Comprehensive Multi-Task Benchmarks (5 papers)
    - [1] Novachart: A large-scale dataset towards chart understanding and generation of multimodal large language models (Linmei Hu, 2024) [View paper](#)
    - [6] ChartMuseum: Testing Visual Reasoning Capabilities of Large Vision-Language Models (Tang Liyan, 2025) [View paper](#)
    - [28] Chartbench: A benchmark for complex visual reasoning in charts (Xu, 2023) [View paper](#)
    - [38] Omni-Chart-600K: A Comprehensive Dataset of Chart Types for Chart Understanding (Shulei Wang, 2025) [View paper](#)
    - [44] ChartBench: A Comprehensive Evaluation Benchmark for Chart Understanding Capabilities of Vision-Language Models (H Long, 2025) [View paper](#)
  - Specialized Benchmarks for Chart Subtypes (4 papers)
    - [26] MultiChartQA: Benchmarking vision-language models on multi-chart problems (Zifeng Zhu, 2025) [View paper](#)
    - [29] Flowlearn: Evaluating large vision-language models on flowchart understanding (Huitong Pan, 2024) [View paper](#)
    - [32] Encqa: Benchmarking vision-language models on visual encodings for charts (Mukherjee, 2025) [View paper](#)
    - [42] FinChart-Bench: Benchmarking Financial Chart Comprehension in Vision-Language Models (Shu Dong, 2025) [View paper](#)
  - Human-Centered and Cognitive Evaluation (2 papers)
    - [12] CHART-6: Human-Centered Evaluation of Data Visualization Understanding in Vision-Language Models (Mukherjee, 2025) [View paper](#)
    - [25] Benchmarking visual language models on standardized visualization literacy tests (Saugat Pandey, 2025) [View paper](#)
  - Robustness and Error Analysis (4 papers)
    - [11] Evaluation and Analysis of Chart Reasoning Accuracy in Multimodal Large Language Models: An Empirical Study on Influencing Factors (Z Jiang, 2025) [View paper](#)
    - [18] Unmasking deceptive visuals: Benchmarking multimodal large language models on misleading chart question answering (Zixin Chen, 2025) [View paper](#)
    - [31] VisGraphVar: A Benchmark Generator for Assessing Variability in Graph Analysis Using Large Vision-Language Models (Camilo Chacón Sartori, 2024) [View paper](#)
    - [37] The Perils of Chart Deception: How Misleading Visualizations Affect Vision-Language Models (Mahbub, 2025) [View paper](#)
  - Hallucination and Factual Consistency Analysis (2 papers)
    - [21] Do llms understand charts? analyzing and correcting factual errors in chart captioning (Chan, 2024) [View paper](#)
    - [46] ChartHal: A Fine-grained Framework Evaluating Hallucination of Large Vision Language Models in Chart Understanding (Wang Xing-qi, 2025) [View paper](#)
  - Meta-Evaluation and Model-as-Judge Frameworks (1 papers)
    - [27] Judging the Judges: Can Large Vision-Language Models Fairly Evaluate Chart Comprehension and Reasoning? (Laskar, 2025) [View paper](#)
  - Data Contamination and Generalization Assessment (1 papers)
    - [43] Learning or Cheating? Assessing Data Contamination in Large Vision-Language Models (Ahmed Masry, 2025) [View paper](#)
- Application Domains and Task-Specific Adaptations
  - Chart Question Answering (4 papers)
    - [15] Chartinsights: Evaluating multimodal large language models for low-level chart question answering (Yifan Wu, 2024) [View paper](#)
    - [16] Advancing multimodal large language models in chart question answering with visualization-referenced instruction tuning (Ye Yilin, 2024) [View paper](#)
    - [17] Chartmind: A comprehensive benchmark for complex real-world multimodal chart question answering (Jingxuan Wei, 2025) [View paper](#)
    - [30] Are large vision language models up to the challenge of chart comprehension and reasoning? an extensive investigation into the capabilities and limitations of llms (Rahman, 2024) [View paper](#)
  - Chart Summarization and Captioning (1 papers)
    - [35] Altchart: Enhancing vlm-based chart summarization through multi-pretex tasks (Omar Moured, 2024) [View paper](#)
  - Visual Grounding and Reference Resolution (1 papers)
    - [20] Refchartqa: Grounding visual answer on chart images through instruction tuning (Vogel, 2025) [View paper](#)
  - Chart-to-Code and Visualization Generation (1 papers)
    - [19] Is gpt-4v (ision) all you need for automating academic data visualization? exploring vision-language models' capability in reproducing academic charts (Ma WeiCheng, 2024) [View paper](#)
  - Chart-to-Table Conversion and Data Extraction (1 papers)
    - [22] Enhancing Financial VQA in Vision Language Models using Intermediate Structured Representations (KUMAR ABHAS, 2025) [View paper](#)
  - Domain-Specific Chart Understanding (3 papers)
    - [3] Patent Figure Classification Using Large Vision-Language Models (Sushil Awale, 2025) [View paper](#)
    - [10] Multimodal retrieval-augmented generation for financial documents: image-centric analysis of charts and tables with large language models (Cheng Jiang, 2025) [View paper](#)
    - [41] Evaluating Compliance with Visualization Guidelines in Diagrams for Scientific Publications Using Large Vision Language Models (Johannes Rückert, 2025) [View paper](#)
- Foundational Studies and Surveys
  - Chart Understanding Surveys and Reviews (2 papers)
    - [5] Rethinking Chart Understanding Using Multimodal Large Language Models (Adrian Tanasă, 2025) [View paper](#)
    - [9] From pixels to insights: A survey on automatic chart understanding in the era of large foundation models (Kung-Hsiang Huang, 2024) [View paper](#)
  - General VLM Surveys and Broader Visual Understanding (3 papers)
    - [23] Do vision-language models really understand visual language? (Hou Yi-fan, 2024) [View paper](#)
    - [36] A Review on Large Language Models for Visual Analytics (Sonbhadra Sanjay Kumar, 2025) [View paper](#)
    - [39] Revisiting Multi-Modal LLM Evaluation (Jian Lu, 2024) [View paper](#)
  - Empirical Studies and Comparative Analyses (4 papers)
    - [13] Charxiv: Charting gaps in realistic chart understanding in multimodal llms (Sanjeev Arora, 2024) [View paper](#)
    - [34] A Comparative Study of CNNs and Vision-Language Models for Chart Image Classification (Bruno Cme, 2025) [View paper](#)

- [40] MMDocBench: Benchmarking Large Vision-Language Models for Fine-Grained Visual Document Understanding (Zhu Fengbin, 2024) [View paper](#)
- [45] Seeing Sequences like Humans: Pattern Classification Driven Time-Series Forecasting via Vision Language Models (Xing-yu Liu, 2025) [View paper](#)

## Narrative

Core task: deep research capabilities in chart analysis using large vision-language models. The field has evolved into several interconnected branches that collectively address how vision-language models can interpret, reason about, and generate insights from visual data representations. Model Architecture and Training Methodologies explores foundational design choices—ranging from specialized encoders like Vary Vision Vocabulary[7] to compact architectures such as Tinychart[8]—and post-training optimization strategies that refine model behavior through reinforcement learning and instruction tuning. Reasoning and Inference Mechanisms investigates how models perform multi-step logical operations, handle complex queries, and integrate external knowledge, with works like Socratic Chart[33] exemplifying structured reasoning pipelines. Evaluation and Benchmarking provides systematic testbeds such as ChartBench[28], MultiChartQA[26], and domain-specific suites like FinChart-Bench[42] to measure performance across diverse chart types and question complexities. Application Domains and Task-Specific Adaptations tailors models to specialized contexts—financial documents (Multimodal Financial Documents[10]), accessibility (AltChart[35]), and deception detection (Deceptive Visuals[18])—while Foundational Studies and Surveys synthesize emerging trends and identify research gaps.

Recent activity highlights tensions between generalist and specialist approaches: some efforts pursue broad competence across chart genres (Comprehensive Chart Understanding[49], Omni-Chart-600K[38]), while others optimize for efficiency or niche tasks (TinyChart[14], Patent Figure Classification[3]). Post-training optimization has become a focal point for enhancing reasoning depth and reducing hallucinations, with Chart Deep Research[0] situated in this branch alongside methods that leverage feedback loops (Text2Vis Feedback[4]) and iterative refinement. Compared to works emphasizing architectural novelty (Novachart[1]) or large-scale pretraining datasets (Charxiv[13]), Chart Deep Research[0] prioritizes reinforcement learning strategies to deepen analytical capabilities after initial training. This positions it within a growing cluster that views chart understanding not merely as pattern recognition but as a reasoning-intensive process requiring targeted post-training interventions, contrasting with purely data-driven scaling approaches and aligning with efforts to improve interpretability and factual grounding in model outputs.

## Related Works in Same Category

No sibling papers were found in the same taxonomy leaf. A taxonomy-subtopic-level comparison will be produced instead.

### Taxonomy-Level Summary

The original leaf focuses on post-training enhancement through reinforcement learning and reward-based optimization to improve chart reasoning after initial model training. The sibling subtopics address complementary aspects of model development: efficient deployment through compression, foundational cross-modal alignment during pre-training, and enhanced visual perception through tokenization improvements. Together, these represent distinct phases and aspects of building chart analysis systems with large vision-language models.

**Similarities:** - All subtopics aim to enhance chart understanding capabilities in vision-language models - Each addresses optimization challenges at different stages of the model lifecycle - All are concerned with improving model performance on chart-specific tasks rather than general vision-language understanding

**Differences:** - The original leaf operates in the post-training phase using reward signals, while Pre-training Objectives works during initial training with alignment tasks - Post-training Optimization focuses on behavioral refinement through RL/policy methods, whereas Efficient Model Design emphasizes architectural efficiency and compression - Vision Vocabulary Enhancement targets the input encoding layer specifically, while the original leaf addresses end-to-end reasoning improvement - The original leaf uses feedback-driven iterative improvement, while Pre-training Objectives establishes foundational cross-modal representations - Efficient Model Design prioritizes deployment constraints (parameters, tokens, compute), while the original leaf prioritizes reasoning quality regardless of model size

**Suggested Search Directions:** - Investigate whether post-training RL methods can be combined with efficient architectures to maintain reasoning quality in compressed models - Explore how pre-training alignment objectives influence the effectiveness of subsequent reinforcement learning for chart tasks - Examine whether enhanced vision vocabularies reduce the need for extensive post-training optimization by improving initial representations

### Sibling Subtopics

- **Efficient Model Design and Compression** (leaves: 1, papers: 2)
  - Scope: Focuses on parameter-efficient architectures, token merging, and computational optimization for resource-constrained deployment.
  - Exclude: Excludes full-scale model training and reasoning frameworks; those belong under Pre-training Objectives or Reasoning Mechanisms.
- **Pre-training Objectives and Cross-modal Alignment** (leaves: 1, papers: 2)
  - Scope: Addresses novel pre-training tasks designed to align chart images with structured data representations or textual descriptions.
  - Exclude: Excludes post-training optimization and reasoning-specific training; those belong under Post-training Optimization or Reasoning Mechanisms.
- **Vision Vocabulary and Encoding Enhancement** (leaves: 1, papers: 1)
  - Scope: Focuses on expanding or optimizing visual tokenization and encoding mechanisms to improve chart perception.
  - Exclude: Excludes general pre-training objectives and reasoning frameworks; those belong under Pre-training Objectives or Reasoning Mechanisms.

## Contributions Analysis

**Overall novelty summary.** The paper proposes PRPO (Parallel Relative Policy Optimization) for post-training optimization and introduces MCDR-Bench to evaluate deep research capabilities in chart analysis. It occupies the 'Post-training Optimization and Reinforcement Learning' leaf within the taxonomy, which currently contains only this work among 50 surveyed papers. This isolation suggests the leaf represents an emerging or underexplored direction: while the broader 'Model Architecture and Training Methodologies' branch includes pre-training strategies and efficient design, dedicated post-training reinforcement learning for chart reasoning appears sparse in the surveyed literature.

The taxonomy reveals neighboring leaves focused on pre-training alignment (e.g., Novachart, Charxiv) and efficient architectures (TinyChart, Vary Vision Vocabulary), but these exclude post-training optimization by design. The 'Reasoning and Inference Mechanisms' branch addresses chain-of-thought and code-driven reasoning yet excludes training methodologies. PRPO's emphasis on disentangling multi-dimensional reward signals and heterogeneous data gradients positions it at the intersection of training innovation and reasoning enhancement, bridging gaps between architectural design and inference-time strategies without directly overlapping either domain.

Among three contributions analyzed, the MCDR-Bench benchmark examined one candidate paper with no refutations found, while PRPO and the unified framework examined zero candidates each. The limited search scope—one candidate total across all contributions—

means the analysis cannot confirm whether substantial prior work exists in parallel policy optimization or deep research evaluation for charts. The absence of refutations reflects the narrow search rather than definitive novelty; a broader literature review covering reinforcement learning in vision-language models or multi-task reward optimization could reveal relevant precedents.

Based on top-K semantic search examining one candidate, the work appears positioned in a sparsely populated taxonomy leaf with minimal direct competition in the surveyed set. However, the restricted search scope leaves open whether related post-training methods in adjacent fields (e.g., general LLM alignment, multi-objective RL) address similar challenges. The taxonomy structure suggests the contribution targets a recognized gap, but comprehensive novelty assessment would require examining reinforcement learning literature beyond chart-specific applications.

---

This paper presents **3 main contributions**, each analyzed against relevant prior work:

### Contribution 1: Parallel Relative Policy Optimization (PRPO) training method

**Description:** PRPO is a training methodology that addresses multi-dimensional reward interference and heterogeneous data gradient conflicts by performing parallel optimization across reward dimensions and partitioning capabilities across data types. This approach enables coordinated development of complex analytical capabilities required for chart deep research.

This contribution was assessed against **0 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

### Contribution 2: MCDR-Bench evaluation benchmark

**Description:** MCDR-Bench is an evaluation benchmark that transforms subjective deep research assessment into objective error identification using the error uniqueness principle. It enables systematic and quantifiable measurement of chart deep research capabilities through controlled error injection.

This contribution was assessed against **1 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

## 1. Overview of pan 2024: multi-author writing style analysis, multilingual text detoxification, oppositional thinking analysis, and generative ai authorship verification

URL: [View paper](#)

### Brief Assessment

PAN 2024[76] focuses on authorship verification and text detoxification tasks, not chart analysis or error injection for evaluating deep research capabilities. The candidate's error injection is for making generated text more human-like, not for objective evaluation of analytical reasoning.

---

### Contribution 3: Unified framework for chart deep research advancement

**Description:** The authors present a unified framework that jointly addresses training and evaluation bottlenecks in chart deep research. By combining PRPO for training and MCDR-Bench for evaluation, the framework provides a systematic approach to developing and measuring advanced analytical reasoning capabilities.

This contribution was assessed against **0 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

## Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

---

## References

- [0] Chart Deep Research in LLMs via Parallel Relative Policy Optimization [View paper](#)
- [1] Novachart: A large-scale dataset towards chart understanding and generation of multimodal large language models [View paper](#)
- [2] End-to-end chart summarization via visual chain-of-thought in vision-language models [View paper](#)
- [3] Patent Figure Classification Using Large Vision-Language Models [View paper](#)
- [4] Closing the feedback loop in text2vis: Refining visualization with vision-language models [View paper](#)
- [5] Rethinking Chart Understanding Using Multimodal Large Language Models [View paper](#)
- [6] ChartMuseum: Testing Visual Reasoning Capabilities of Large Vision-Language Models [View paper](#)
- [7] Vary: Scaling up the Vision Vocabulary for Large Vision-Language Models [View paper](#)
- [8] Tinychart: Efficient chart understanding with visual token merging and program-of-thoughts learning [View paper](#)
- [9] From pixels to insights: A survey on automatic chart understanding in the era of large foundation models [View paper](#)
- [10] Multimodal retrieval-augmented generation for financial documents: image-centric analysis of charts and tables with large language models [View paper](#)
- [11] Evaluation and Analysis of Chart Reasoning Accuracy in Multimodal Large Language Models: An Empirical Study on Influencing Factors [View paper](#)
- [12] CHART-6: Human-Centered Evaluation of Data Visualization Understanding in Vision-Language Models [View paper](#)
- [13] Charxiv: Charting gaps in realistic chart understanding in multimodal llms [View paper](#)
- [14] TinyChart: Efficient Chart Understanding with Program-of-Thoughts Learning and Visual Token Merging [View paper](#)
- [15] Chartinsights: Evaluating multimodal large language models for low-level chart question answering [View paper](#)
- [16] Advancing multimodal large language models in chart question answering with visualization-referenced instruction tuning [View paper](#)
- [17] Chartmind: A comprehensive benchmark for complex real-world multimodal chart question answering [View paper](#)
- [18] Unmasking deceptive visuals: Benchmarking multimodal large language models on misleading chart question answering [View paper](#)
- [19] Is gpt-4v (ision) all you need for automating academic data visualization? exploring vision-language models' capability in reproducing academic charts [View paper](#)
- [20] Refchartqa: Grounding visual answer on chart images through instruction tuning [View paper](#)
- [21] Do llms understand charts? analyzing and correcting factual errors in chart captioning [View paper](#)
- [22] Enhancing Financial VQA in Vision Language Models using Intermediate Structured Representations [View paper](#)
- [23] Do vision-language models really understand visual language? [View paper](#)
- [24] ChartSketcher: Reasoning with Multimodal Feedback and Reflection for Chart Understanding [View paper](#)
- [25] Benchmarking visual language models on standardized visualization literacy tests [View paper](#)
- [26] MultiChartQA: Benchmarking vision-language models on multi-chart problems [View paper](#)

- [27] Judging the Judges: Can Large Vision-Language Models Fairly Evaluate Chart Comprehension and Reasoning? [View paper](#)
- [28] ChartBench: A benchmark for complex visual reasoning in charts [View paper](#)
- [29] Flowlearn: Evaluating large vision-language models on flowchart understanding [View paper](#)
- [30] Are large vision language models up to the challenge of chart comprehension and reasoning? an extensive investigation into the capabilities and limitations of vlms [View paper](#)
- [31] VisGraphVar: A Benchmark Generator for Assessing Variability in Graph Analysis Using Large Vision-Language Models [View paper](#)
- [32] Encqa: Benchmarking vision-language models on visual encodings for charts [View paper](#)
- [33] Socratic chart: Cooperating multiple agents for robust svg chart understanding [View paper](#)
- [34] A Comparative Study of CNNs and Vision-Language Models for Chart Image Classification [View paper](#)
- [35] Altchart: Enhancing vlm-based chart summarization through multi-pretext tasks [View paper](#)
- [36] A Review on Large Language Models for Visual Analytics [View paper](#)
- [37] The Perils of Chart Deception: How Misleading Visualizations Affect Vision-Language Models [View paper](#)
- [38] Omni-Chart-600K: A Comprehensive Dataset of Chart Types for Chart Understanding [View paper](#)
- [39] Revisiting Multi-Modal LLM Evaluation [View paper](#)
- [40] MMDocBench: Benchmarking Large Vision-Language Models for Fine-Grained Visual Document Understanding [View paper](#)
- [41] Evaluating Compliance with Visualization Guidelines in Diagrams for Scientific Publications Using Large Vision Language Models [View paper](#)
- [42] FinChart-Bench: Benchmarking Financial Chart Comprehension in Vision-Language Models [View paper](#)
- [43] Learning or Cheating? Assessing Data Contamination in Large Vision-Language Models [View paper](#)
- [44] ChartBench: A Comprehensive Evaluation Benchmark for Chart Understanding Capabilities of Vision-Language Models [View paper](#)
- [45] Seeing Sequences like Humans: Pattern Classification Driven Time-Series Forecasting via Vision Language Models [View paper](#)
- [46] ChartHal: A Fine-grained Framework Evaluating Hallucination of Large Vision Language Models in Chart Understanding [View paper](#)
- [47] ChartGen: Scaling Chart Understanding Via Code-Guided Synthetic Chart Generation [View paper](#)
- [48] Chart-CoCa: Self-Improving Chart Understanding of Vision LMs via Code-Driven Synthesis and Candidate-Conditioned Answering [View paper](#)
- [49] In-Depth and In-Breadth: Pre-training Multimodal Language Models Customized for Comprehensive Chart Understanding [View paper](#)
- [50] Enhanced Chart Understanding in Vision and Language Task via Cross-modal Pre-training on Plot Table Pairs [View paper](#)
- [51] Agentthink: A unified framework for tool-augmented chain-of-thought reasoning in vision-language models for autonomous driving [View paper](#)
- [52] Effectivity of interactive multimedia with theocentric approach to the analytical thinking skills of elementary school students in science learning [View paper](#)
- [53] Nature's Insight: A Novel Framework and Comprehensive Analysis of Agentic Reasoning Through the Lens of Neuroscience [View paper](#)
- [54] Mapping the Minds of LLMs: A Graph-Based Analysis of Reasoning LLM [View paper](#)
- [55] Contrastive learning framework for multimodal knowledge graph construction and data-analytical reasoning [View paper](#)
- [56] Learning to align multi-faceted evaluation: A unified and robust framework [View paper](#)
- [57] Time-MQA: Time Series Multi-Task Question Answering with Context Enhancement [View paper](#)
- [58] From Perception to Cognition: A Survey of Vision-Language Interactive Reasoning in Multimodal Large Language Models [View paper](#)
- [59] GRAFT: GRaPH and Table Reasoning for Textual Alignment - A Benchmark for Structured Instruction Following and Visual Reasoning [View paper](#)
- [60] The impact of generative AI on critical thinking skills: a systematic review, conceptual framework and future research directions [View paper](#)
- [61] Scalable parallel task scheduling for autonomous driving using multi-task deep reinforcement learning [View paper](#)
- [62] Dynamic weights in multi-objective deep reinforcement learning [View paper](#)
- [63] Towards Optimized Dynamic Ridesharing System Through Multi-Objective Reinforcement Learning [View paper](#)
- [64] Safe and balanced: A framework for constrained multi-objective reinforcement learning [View paper](#)
- [65] Welfare and fairness in multi-objective reinforcement learning [View paper](#)
- [66] Rewards-in-context: Multi-objective alignment of foundation models with dynamic preference adjustment [View paper](#)
- [67] Preference-Based Multi-Objective Reinforcement Learning [View paper](#)
- [68] Distributional pareto-optimal multi-objective reinforcement learning [View paper](#)
- [69] Beyond Accuracy: Decision Transformers for Reward-Driven Multi-Objective Recommendations [View paper](#)
- [70] REMOR: Automated Peer Review Generation with LLM Reasoning and Multi-Objective Reinforcement Learning [View paper](#)
- [71] Ratescore: A metric for radiology report generation [View paper](#)
- [72] AudioLDM: Text-to-Audio Generation with Latent Diffusion Models [View paper](#)
- [73] TIGERScore: Towards Building Explainable Metric for All Text Generation Tasks [View paper](#)
- [74] MQM-APE: toward high-quality error annotation predictors with automatic post-editing in LLM translation evaluators [View paper](#)
- [75] On the Blind Spots of Model-Based Evaluation Metrics for Text Generation [View paper](#)
- [76] Overview of pan 2024: multi-author writing style analysis, multilingual text detoxification, oppositional thinking analysis, and generative ai authorship verification [View paper](#)
- [77] Speechbertscore: Reference-aware automatic evaluation of speech generation leveraging nlp evaluation metrics [View paper](#)
- [78] Towards explainable evaluation metrics for natural language generation [View paper](#)
- [79] Instructscore: Explainable text generation evaluation with finegrained feedback [View paper](#)
- [80] Electric load forecasting under false data injection attacks via denoising deep learning and generative adversarial networks [View paper](#)